Oracle® Database

Advanced Application Developer's Guide

11*g* Release 2 (11.2)

E17125-04

October 2010

ORACLE®

Oracle Database Advanced Application Developer's Guide, 11*g* Release 2 (11.2)

E17125-04

Copyright © 1996, 2010, Oracle and/or its affiliates. All rights reserved.

Primary Author: Sheila Moore

Contributing Authors: D. Adams, L. Ashdown, M. Cowan, J. Melnick, R. Moran, E. Paapanen, J. Russell, R. Strohm, R. Ward

Contributors: D. Alpern, G. Arora, C. Barclay, D. Bronnikov, T. Chang, L. Chen, B. Cheng, M. Davidson, R. Day, R. Decker, G. Doherty, D. Elson, A. Ganesh, M. Hartstein, Y. Hu, J. Huang, C. Iyer, N. Jain, R. Jenkins Jr., S. Kotsovolos, V. Krishnaswamy, S. Kumar, C. Lei, B. Llewellyn, D. Lorentz, V. Moore, K. Muthukkaruppan, V. Moore, J. Muller, R. Murthy, R. Pang, B. Sinha, S. Vemuri, W. Wang, D. Wong, A. Yalamanchi, Q. Yu

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

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

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface xxvii

Audience xxvii
Documentation Accessibility xxvii
Related Documents xxviii
Conventions xxviii

What's New in Application Development? xxxi

Oracle Database 11*g* Release 2 (11.2.0.2) Feature xxxi
Oracle Database 11*g* Release 2 (11.2.0.1) Features xxxi
Oracle Database 11*g* Release 1 (11.1) Features xxxiv

Part I SQL for Application Developers

1 SQL Processing for Application Developers

Description of SQL Statement Processing 1-1
Grouping Operations into Transactions 1-4
Deciding How to Group Operations in Transactions 1-4
Improving Transaction Performance 1-4
Committing Transactions 1-5
Managing Commit Redo Action 1-6
Rolling Back Transactions 1-7
Defining Transaction Savepoints 1-8
Ensuring Repeatable Reads with Read-Only Transactions 1-8
Using Cursors 1-9
How Many Cursors Can a Session Have? 1-9
Using a Cursor to Reexecute a Statement 1-10
Scrollable Cursors 1-10
Closing a Cursor 1-10
Canceling a Cursor 1-11
Locking Tables Explicitly 1-11
Privileges Required to Acquire Table Locks 1-12
Choosing a Locking Strategy 1-12
When to Lock with ROW SHARE MODE and ROW EXCLUSIVE MODE 1-13
When to Lock with SHARE MODE 1-13
When to Lock with SHARE ROW EXCLUSIVE MODE 1-14

When to Lock with EXCLUSIVE MODE 1-14
Letting Oracle Database Control Table Locking 1-14
Explicitly Acquiring Row Locks 1-15
Examples of Concurrency Under Explicit Locking 1-16
Using Oracle Lock Management Services (User Locks) 1-22
When to Use User Locks 1-23
Viewing and Monitoring Locks 1-23
Using Serializable Transactions for Concurrency Control 1-24
Transaction Interaction and Isolation Level 1-25
Setting Isolation Levels 1-26
Serializable Transactions and Referential Integrity 1-27
READ COMMITTED and SERIALIZABLE Isolation Levels 1-28
Transaction Set Consistency Differences 1-29
Choosing Transaction Isolation Levels 1-29
Autonomous Transactions 1-30
Examples of Autonomous Transactions 1-32
Ordering a Product 1-33
Withdrawing Money from a Bank Account 1-33
Defining Autonomous Transactions 1-36
Resuming Execution After Storage Allocation Errors 1-37
What Operations Have Resumable Storage Allocation? 1-37
Handling Suspended Storage Allocation 1-37
Using an AFTER SUSPEND Trigger in the Application 1-38
Checking for Suspended Statements 1-39

2 Using SQL Data Types in Database Applications

Overview of SQL Data Types 2-2
Representing Character Data 2-2
Specifying Column Lengths as Bytes or Characters 2-3
Choosing Between CHAR and VARCHAR2 Data Types 2-4
Representing Numeric Data 2-4
Floating-Point Number Components 2-5
Floating-Point Number Formats 2-5
Binary Floating-Point Formats 2-6
Special Values for Native Floating-Point Formats 2-7
Comparison Operators for Native Floating-Point Data Types 2-8
Arithmetic Operations with Native Floating-Point Data Types 2-8
Conversion Functions for Floating-Point Data Types 2-9
Client Interfaces for Native Floating-Point Data Types 2-10
OCI Native Floating-Point Data Types SQLT_BFLOAT and SQLT_BDOUBLE 2-10
Native Floating-Point Data Types Supported in ADTs 2-10
Pro*C/C++ Support for Native Floating-Point Data Types 2-10
Representing Date and Time Data 2-10
Displaying Current Date and Time 2-11
Changing the Default Date Format 2-12
Changing the Default Time Format 2-12
Arithmetic Operations with Date and Time Data Types 2-13

Converting Between Date and Time Types ... 2-13
Importing and Exporting Date and Time Types ... 2-14
Representing Specialized Data ... 2-14
Representing Geographic Data ... 2-14
Representing Multimedia Data ... 2-14
Representing Large Amounts of Data ... 2-15
Representing Searchable Text ... 2-16
Representing XML ... 2-16
Representing Dynamically Typed Data ... 2-17
Representing Data with ANSI/ISO, DB2, and SQL/DS Data Types ... 2-19
Representing Conditional Expressions as Data ... 2-20
Identifying Rows by Address ... 2-21
Querying the ROWID Pseudocolumn ... 2-22
ROWID Data Type ... 2-23
Restricted ROWID ... 2-23
Extended ROWID ... 2-23
External Binary ROWID ... 2-24
UROWID Data Type ... 2-24
How Oracle Database Converts Data Types ... 2-24
Data Type Conversion During Assignments ... 2-24
Data Type Conversion During Expression Evaluation ... 2-25
Metadata for SQL Built-In Functions ... 2-26

3 Using Regular Expressions in Database Applications

Overview of Regular Expressions ... 3-1
What Are Regular Expressions? ... 3-1
How Are Regular Expressions Useful? ... 3-2
Oracle Database Implementation of Regular Expressions ... 3-2
Oracle Database Support for the POSIX Regular Expression Standard ... 3-4
Metacharacters in Regular Expressions ... 3-4
POSIX Metacharacters in Oracle Database Regular Expressions ... 3-4
Multilingual Extensions to POSIX Regular Expression Standard ... 3-7
PERL-Influenced Extensions to POSIX Regular Expression Standard ... 3-8
Using Regular Expressions in SQL Statements: Scenarios ... 3-10
Using a Constraint to Enforce a Phone Number Format ... 3-10
Using Back References to Reposition Characters ... 3-11

4 Using Indexes in Database Applications

Privileges Needed to Create Indexes ... 4-1
Guidelines for Application-Specific Indexes ... 4-1
Which Come First, Data or Indexes? ... 4-2
Create a Temporary Table Space Before Creating Indexes ... 4-2
Index the Correct Tables and Columns ... 4-3
Limit the Number of Indexes for Each Table ... 4-4
Choose Column Order in Composite Indexes ... 4-4
Gather Index Statistics ... 4-5

Drop Unused Indexes 4-5
Examples of Creating Basic Indexes 4-6
When to Use Domain Indexes 4-7
When to Use Function-Based Indexes 4-7
Advantages of Function-Based Indexes 4-8
Restrictions on Function-Based Indexes 4-10
Examples of Function-Based Indexes 4-12
Function-Based Index for Case-Insensitive Searches 4-12
Precomputing Arithmetic Expressions with a Function-Based Index 4-12
Function-Based Index for Language-Dependent Sorting 4-12

5 Maintaining Data Integrity in Database Applications

Overview of Constraints 5-1
Enforcing Business Rules with Constraints 5-2
Enforcing Business Rules with Application Logic 5-2
Creating Indexes for Use with Constraints 5-2
When to Use NOT NULL Constraints 5-3
When to Use Default Column Values 5-4
Setting Default Column Values 5-4
Choosing a Primary Key for a Table 5-5
When to Use UNIQUE Constraints 5-5
When to Use Constraints On Views 5-6
Enforcing Referential Integrity with Constraints 5-6
FOREIGN KEY Constraints and NULL Values 5-8
Defining Relationships Between Parent and Child Tables 5-9
Rules for Multiple FOREIGN KEY Constraints 5-10
Deferring Constraint Checks 5-10
Minimizing Space and Time Overhead for Indexes Associated with Constraints 5-12
Guidelines for Indexing Foreign Keys 5-12
Referential Integrity in a Distributed Database 5-12
When to Use CHECK Constraints 5-13
Restrictions on CHECK Constraints 5-13
Designing CHECK Constraints 5-14
Rules for Multiple CHECK Constraints 5-14
Choosing Between CHECK and NOT NULL Constraints 5-14
Examples of Defining Constraints 5-15
Privileges Needed to Define Constraints 5-16
Naming Constraints 5-16
Enabling and Disabling Constraints 5-16
Why Disable Constraints? 5-17
Creating Enabled Constraints (Default) 5-17
Creating Disabled Constraints 5-18
Enabling Existing Constraints 5-18
Disabling Existing Constraints 5-19
Guidelines for Enabling and Disabling Key Constraints 5-19
Fixing Constraint Exceptions 5-19
Modifying Constraints 5-20

Renaming Constraints ... 5-21
Dropping Constraints ... 5-22
Managing FOREIGN KEY Constraints ... 5-22
Data Types and Names for Foreign Key Columns ... 5-22
Limit on Columns in Composite Foreign Keys ... 5-23
Foreign Key References Primary Key by Default ... 5-23
Privileges Required to Create FOREIGN KEY Constraints ... 5-23
Choosing How Foreign Keys Enforce Referential Integrity ... 5-23
Viewing Information About Constraints ... 5-24

Part II PL/SQL for Application Developers

6 Coding PL/SQL Subprograms and Packages

Overview of PL/SQL Units ... 6-1
Anonymous Blocks ... 6-2
Stored PL/SQL Units ... 6-4
Naming Subprograms ... 6-5
Subprogram Parameters ... 6-5
Creating Subprograms ... 6-8
Altering Subprograms ... 6-9
Dropping Subprograms and Packages ... 6-10
External Subprograms ... 6-10
PL/SQL Function Result Cache ... 6-10
PL/SQL Packages ... 6-10
PL/SQL Object Size Limits ... 6-13
Creating Packages ... 6-13
Naming Packages and Package Objects ... 6-14
Package Invalidations and Session State ... 6-14
Packages Supplied with Oracle Database ... 6-15
Overview of Bulk Binding ... 6-16
When to Use Bulk Binds ... 6-16
Triggers ... 6-18
Compiling PL/SQL Subprograms for Native Execution ... 6-18
Cursor Variables ... 6-19
Declaring and Opening Cursor Variables ... 6-19
Examples of Cursor Variables ... 6-20
Handling PL/SQL Compile-Time Errors ... 6-22
Handling Run-Time PL/SQL Errors ... 6-23
Declaring Exceptions and Exception Handlers ... 6-24
Unhandled Exceptions ... 6-25
Handling Errors in Distributed Queries ... 6-25
Handling Errors in Remote Subprograms ... 6-26
Debugging Stored Subprograms ... 6-26
PL/Scope ... 6-27
PL/SQL Hierarchical Profiler ... 6-27
Oracle JDeveloper ... 6-27

DBMS_OUTPUT Package 6-27
Privileges for Debugging PL/SQL and Java Stored Subprograms 6-27
Writing Low-Level Debugging Code 6-28
DBMS_DEBUG_JDWP Package 6-29
DBMS_DEBUG Package 6-29
Invoking Stored Subprograms 6-29
Privileges Required to Invoke a Subprogram 6-30
Invoking a Subprogram Interactively from Oracle Tools 6-30
Invoking a Subprogram from Another Subprogram 6-32
Invoking a Subprogram from a 3GL Application 6-33
Invoking Remote Subprograms 6-33
Synonyms for Remote Subprograms 6-34
Committing Transactions 6-35
Invoking Stored PL/SQL Functions from SQL Statements 6-35
Why Invoke Stored PL/SQL Subprograms from SQL Statements? 6-36
Where PL/SQL Functions Can Appear in SQL Statements 6-36
When PL/SQL Functions Can Appear in SQL Expressions 6-37
Controlling Side Effects 6-38
Restrictions 6-39
Declaring a Function 6-39
Parallel Query and Parallel DML 6-40
PRAGMA RESTRICT_REFERENCES for Backward Compatibility 6-41
Returning Large Amounts of Data from a Function 6-45
Coding Your Own Aggregate Functions 6-45

7 Using PL/Scope

Specifying Identifier Collection 7-1
PL/Scope Identifier Data for STANDARD and DBMS_STANDARD 7-2
How Much Space is PL/Scope Data Using? 7-4
Viewing PL/Scope Data 7-5
Static Data Dictionary Views 7-5
Unique Keys 7-5
Context 7-5
Signature 7-7
Demo Tool 7-7
SQL Developer 7-8
Identifier Types that PL/Scope Collects 7-8
Usages that PL/Scope Reports 7-9
Sample PL/Scope Session 7-10

8 Using the PL/SQL Hierarchical Profiler

Overview of PL/SQL Hierarchical Profiler 8-1
Collecting Profile Data 8-2
Understanding Raw Profiler Output 8-3
Namespaces of Tracked Subprograms 8-6
Special Function Names 8-6
Analyzing Profile Data 8-6

Creating Hierarchical Profiler Tables 8-7
Understanding Hierarchical Profiler Tables 8-8
Hierarchical Profiler Database Table Columns 8-8
Distinguishing Between Overloaded Subprograms 8-10
Hierarchical Profiler Tables for Sample PL/SQL Procedure 8-10
Examples of Calls to DBMS_HPROF.analyze with Options 8-11
plshprof Utility 8-13
plshprof Options 8-13
HTML Report from a Single Raw Profiler Output File 8-14
First Page of Report 8-14
Function-Level Reports 8-15
Module-Level Reports 8-16
Namespace-Level Reports 8-16
Parents and Children Report for a Function 8-17
HTML Difference Report from Two Raw Profiler Output Files 8-18
Difference Report Conventions 8-19
First Page of Difference Report 8-19
Function-Level Difference Reports 8-20
Module-Level Difference Reports 8-21
Namespace-Level Difference Reports 8-22
Parents and Children Difference Report for a Function 8-22

9 Developing PL/SQL Web Applications

Overview of PL/SQL Web Applications 9-1
Implementing PL/SQL Web Applications 9-2
PL/SQL Gateway 9-2
mod_plsql 9-2
Embedded PL/SQL Gateway 9-3
PL/SQL Web Toolkit 9-3
Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application 9-4
Using Embedded PL/SQL Gateway 9-4
How Embedded PL/SQL Gateway Processes Client Requests 9-5
Installing Embedded PL/SQL Gateway 9-6
Configuring Embedded PL/SQL Gateway 9-6
Configuring Embedded PL/SQL Gateway: Overview 9-6
Configuring User Authentication for Embedded PL/SQL Gateway 9-8
Invoking PL/SQL Stored Subprograms Through Embedded PL/SQL Gateway 9-17
Securing Application Access with Embedded PL/SQL Gateway 9-17
Restrictions in Embedded PL/SQL Gateway 9-18
Using Embedded PL/SQL Gateway: Scenario 9-18
Generating HTML Output with PL/SQL 9-20
Passing Parameters to PL/SQL Web Applications 9-21
Passing List and Dropdown-List Parameters from an HTML Form 9-21
Passing Option and Check Box Parameters from an HTML Form 9-22
Passing Entry-Field Parameters from an HTML Form 9-22
Passing Hidden Parameters from an HTML Form 9-24
Uploading a File from an HTML Form 9-24

Submitting a Completed HTML Form 9-24
Handling Missing Input from an HTML Form 9-25
Maintaining State Information Between Web Pages 9-25
Performing Network Operations in PL/SQL Subprograms 9-25
Sending E-Mail from PL/SQL 9-26
Getting a Host Name or Address from PL/SQL 9-27
Using TCP/IP Connections from PL/SQL 9-27
Retrieving HTTP URL Contents from PL/SQL 9-27
Using Tables, Image Maps, Cookies, and CGI Variables from PL/SQL 9-29

10 Developing PL/SQL Server Pages (PSP)

What Are PL/SQL Server Pages and Why Use Them? 10-1
Prerequisites for Developing and Deploying PL/SQL Server Pages 10-2
PL/SQL Server Pages and the HTP Package 10-3
PL/SQL Server Pages and Other Scripting Solutions 10-3
Developing PL/SQL Server Pages 10-4
Specifying Basic Server Page Characteristics 10-5
Specifying the Scripting Language 10-6
Returning Data to the Client Browser 10-6
Handling Script Errors 10-7
Accepting User Input 10-8
Naming the PL/SQL Stored Procedure 10-9
Including the Contents of Other Files 10-9
Declaring Global Variables in a PSP Script 10-10
Specifying Executable Statements in a PSP Script 10-10
Substituting Expression Values in a PSP Script 10-11
Using Quotation Marks and Escaping Strings in a PSP Script 10-12
Including Comments in a PSP Script 10-12
Loading PL/SQL Server Pages into the Database 10-13
Querying PL/SQL Server Page Source Code 10-14
Running PL/SQL Server Pages Through URLs 10-15
Examples of PL/SQL Server Pages 10-16
Setup for PL/SQL Server Pages Examples 10-16
Printing the Sample Table with a Loop 10-17
Allowing a User Selection 10-18
Using an HTML Form to Invoke a PL/SQL Server Page 10-19
Including JavaScript in a PSP File 10-20
Debugging PL/SQL Server Pages 10-21
Putting PL/SQL Server Pages into Production 10-22

11 Using Continuous Query Notification (CQN)

Object Change Notification (OCN) 11-2
Query Result Change Notification (QRCN) 11-2
Guaranteed Mode 11-3
Best-Effort Mode 11-3
Events that Generate Notifications 11-5
Committed DML Transactions 11-5

Committed DDL Statements 11-5
Deregistration 11-6
Global Events 11-6
Notification Contents 11-7
Good Candidates for CQN 11-7
Creating CQN Registrations 11-10
PL/SQL CQN Registration Interface 11-10
CQN Registration Options 11-11
Notification Type Option 11-11
QRCN Mode (QRCN Notification Type Only) 11-11
ROWID Option 11-12
Operations Filter Option (OCN Notification Type Only) 11-12
Transaction Lag Option (OCN Notification Type Only) 11-13
Notification Grouping Options 11-13
Reliable Option 11-14
Purge-on-Notify and Timeout Options 11-14
Prerequisites for Creating CQN Registrations 11-14
Queries that Can Be Registered for Object Change Notification (OCN) 11-15
Queries that Can Be Registered for Query Result Change Notification (QRCN) 11-15
Queries that Can Be Registered for QRCN in Guaranteed Mode 11-15
Queries that Can Be Registered for QRCN Only in Best-Effort Mode 11-16
Queries that Cannot Be Registered for QRCN in Either Mode 11-17
Using PL/SQL to Register Queries for CQN 11-18
Creating a PL/SQL Notification Handler 11-18
Creating a CQ_NOTIFICATION$_REG_INFO Object 11-19
Identifying Individual Queries in a Notification 11-22
Adding Queries to an Existing Registration 11-22
Best Practices for CQN Registrations 11-22
Troubleshooting CQN Registrations 11-23
Querying CQN Registrations 11-24
Interpreting Notifications 11-24
Interpreting a CQ_NOTIFICATION$_DESCRIPTOR Object 11-24
Interpreting a CQ_NOTIFICATION$_TABLE Object 11-25
Interpreting a CQ_NOTIFICATION$_QUERY Object 11-26
Interpreting a CQ_NOTIFICATION$_ROW Object 11-26
Deleting Registrations 11-27
Configuring CQN: Scenario 11-27
Creating a PL/SQL Notification Handler 11-27
Registering the Queries 11-29

Part III Advanced Topics for Application Developers

12 Using Oracle Flashback Technology

Overview of Oracle Flashback Technology 12-1
Application Development Features 12-2
Database Administration Features 12-3

Configuring Your Database for Oracle Flashback Technology ... 12-3
Configuring Your Database for Automatic Undo Management ... 12-4
Configuring Your Database for Oracle Flashback Transaction Query ... 12-4
Configuring Your Database for Flashback Transaction ... 12-4
Enabling Oracle Flashback Operations on Specific LOB Columns ... 12-5
Granting Necessary Privileges ... 12-5
Using Oracle Flashback Query (SELECT AS OF) ... 12-6
Example of Examining and Restoring Past Data ... 12-7
Guidelines for Oracle Flashback Query ... 12-7
Using Oracle Flashback Version Query ... 12-8
Using Oracle Flashback Transaction Query ... 12-9
Using Oracle Flashback Transaction Query with Oracle Flashback Version Query ... 12-10
Using ORA_ROWSCN ... 12-12
Scenario: Package Subprogram Might Change Row ... 12-13
ORA_ROWSCN and Tables with Virtual Private Database (VPD) ... 12-13
Using DBMS_FLASHBACK Package ... 12-14
Using Flashback Transaction ... 12-15
Dependent Transactions ... 12-16
TRANSACTION_BACKOUT Parameters ... 12-16
TRANSACTION_BACKOUT Reports ... 12-17
*_FLASHBACK_TXN_STATE ... 12-17
*_FLASHBACK_TXN_REPORT ... 12-17
Using Flashback Data Archive (Oracle Total Recall) ... 12-18
Creating a Flashback Data Archive ... 12-18
Altering a Flashback Data Archive ... 12-19
Dropping a Flashback Data Archive ... 12-20
Specifying the Default Flashback Data Archive ... 12-20
Enabling and Disabling Flashback Data Archive ... 12-21
DDL Statements on Tables Enabled for Flashback Data Archive ... 12-22
Viewing Flashback Data Archive Data ... 12-22
Flashback Data Archive Scenarios ... 12-23
Scenario: Using Flashback Data Archive to Enforce Digital Shredding ... 12-23
Scenario: Using Flashback Data Archive to Access Historical Data ... 12-23
Scenario: Using Flashback Data Archive to Generate Reports ... 12-24
Scenario: Using Flashback Data Archive for Auditing ... 12-24
Scenario: Using Flashback Data Archive to Recover Data ... 12-25
General Guidelines for Oracle Flashback Technology ... 12-25
Performance Guidelines for Oracle Flashback Technology ... 12-26

13 Choosing a Programming Environment

Overview of Application Architecture ... 13-2
Client/Server Architecture ... 13-2
Server-Side Programming ... 13-2
Two-Tier and Three-Tier Architecture ... 13-2
Overview of the Program Interface ... 13-3
User Interface ... 13-3
Stateful and Stateless User Interfaces ... 13-3

Overview of PL/SQL 13-4
Overview of Oracle Database Java Support 13-4
Overview of Oracle JVM 13-5
Overview of Oracle JDBC 13-5
Oracle JDBC Drivers 13-6
Sample JDBC 2.0 Program 13-7
Sample Pre-2.0 JDBC Program 13-8
Overview of Oracle SQLJ 13-8
Benefits of SQLJ 13-9
SQLJ Stored Subprograms in the Server 13-10
Comparing Oracle JDBC and Oracle SQLJ 13-10
Overview of Oracle JPublisher 13-11
Overview of Java Stored Subprograms 13-11
Overview of Oracle Database Web Services 13-12
Choosing PL/SQL or Java 13-13
Similarities of PL/SQL and Java 13-13
PL/SQL Advantages Over Java 13-14
Java Advantages Over PL/SQL 13-14
Overview of Precompilers 13-14
Overview of the Pro*C/C++ Precompiler 13-14
Overview of the Pro*COBOL Precompiler 13-16
Overview of OCI and OCCI 13-18
Advantages of OCI and OCCI 13-19
OCI and OCCI Functions 13-19
Procedural and Nonprocedural Elements of OCI and OCCI Applications 13-19
Building an OCI or OCCI Application 13-20
Choosing a Precompiler or OCI 13-21
Overview of Oracle Data Provider for .NET (ODP.NET) 13-21
Overview of OraOLEDB 13-22
Overview of Oracle Objects for OLE (OO4O) 13-22
OO4O Automation Server 13-23
OO4O Object Model 13-24
OraSession 13-25
OraServer 13-25
OraDatabase 13-25
OraDynaset 13-26
OraField 13-26
OraMetaData and OraMDAttribute 13-26
OraParameter and OraParameters 13-26
OraParamArray 13-27
OraSQLStmt 13-27
OraAQ 13-27
OraAQMsg 13-27
OraAQAgent 13-27
Support for Oracle LOB and Object Data Types 13-28
OraBLOB and OraCLOB 13-28
OraBFILE 13-28

Oracle Data Control 13-29
Oracle Objects for OLE C++ Class Library 13-29

14 Developing Applications with Multiple Programming Languages

Overview of Multilanguage Programs 14-1
What Is an External Procedure? 14-2
Overview of Call Specification for External Procedures 14-3
Loading External Procedures 14-4
Loading Java Class Methods 14-4
Loading External C Procedures 14-4
Define the C Procedures 14-5
Set Up the Environment 14-6
Identify the DLL 14-7
Publish the External Procedures 14-9
Publishing External Procedures 14-9
AS LANGUAGE Clause for Java Class Methods 14-10
AS LANGUAGE Clause for External C Procedures 14-10
LIBRARY 14-10
NAME 14-10
LANGUAGE 14-11
CALLING STANDARD 14-11
WITH CONTEXT 14-11
PARAMETERS 14-11
AGENT IN 14-11
Publishing Java Class Methods 14-11
Publishing External C Procedures 14-12
Locations of Call Specifications 14-12
Example: Locating a Call Specification in a PL/SQL Package 14-13
Example: Locating a Call Specification in a PL/SQL Package Body 14-13
Example: Locating a Call Specification in an ADT Specification 14-13
Example: Locating a Call Specification in an ADT Body 14-14
Example: Java with AUTHID 14-14
Example: C with Optional AUTHID 14-14
Example: Mixing Call Specifications in a Package 14-14
Passing Parameters to External C Procedures with Call Specifications 14-15
Specifying Data Types 14-16
External Data Type Mappings 14-17
Passing Parameters BY VALUE or BY REFERENCE 14-19
Declaring Formal Parameters 14-19
Overriding Default Data Type Mapping 14-20
Specifying Properties 14-20
INDICATOR 14-22
LENGTH and MAXLEN 14-22
CHARSETID and CHARSETFORM 14-22
Repositioning Parameters 14-23
SELF 14-23
BY REFERENCE 14-25

WITH CONTEXT 14-26
Interlanguage Parameter Mode Mappings 14-26
Running External Procedures with CALL Statements 14-26
Preconditions for External Procedures 14-27
Privileges of External Procedures 14-27
Managing Permissions 14-28
Creating Synonyms for External Procedures 14-28
CALL Statement Syntax 14-28
Calling Java Class Methods 14-29
Calling External C Procedures 14-29
Handling Errors and Exceptions in Multilanguage Programs 14-30
Using Service Routines with External C Procedures 14-30
OCIExtProcAllocCallMemory 14-30
OCIExtProcRaiseExcp 14-34
OCIExtProcRaiseExcpWithMsg 14-35
Doing Callbacks with External C Procedures 14-36
OCIExtProcGetEnv 14-36
Object Support for OCI Callbacks 14-38
Restrictions on Callbacks 14-38
Debugging External Procedures 14-39
Example: Calling an External Procedure 14-39
Global Variables in External C Procedures 14-40
Static Variables in External C Procedures 14-40
Restrictions on External C Procedures 14-41

15 Developing Applications with Oracle XA

X/Open Distributed Transaction Processing (DTP) 15-1
DTP Terminology 15-2
Required Public Information 15-4
Oracle XA Library Subprograms 15-5
Oracle XA Library Subprograms 15-5
Oracle XA Interface Extensions 15-6
Developing and Installing XA Applications 15-6
DBA or System Administrator Responsibilities 15-7
Application Developer Responsibilities 15-8
Defining the xa_open String 15-8
Syntax of the xa_open String 15-8
Required Fields for the xa_open String 15-9
Optional Fields for the xa_open String 15-9
Using Oracle XA with Precompilers 15-11
Using Precompilers with the Default Database 15-11
Using Precompilers with a Named Database 15-11
Using Oracle XA with OCI 15-12
Managing Transaction Control with Oracle XA 15-13
Examples of Precompiler Applications 15-13
Migrating Precompiler or OCI Applications to TPM Applications 15-14
Managing Oracle XA Library Thread Safety 15-15

Specifying Threading in the Open String 15-16
Restrictions on Threading in Oracle XA 15-16
Using the DBMS_XA Package 15-16
Troubleshooting XA Applications 15-19
Accessing Oracle XA Trace Files 15-19
xa_open String DbgFl 15-20
Trace File Locations 15-20
Managing In-Doubt or Pending Oracle XA Transactions 15-20
Using SYS Account Tables to Monitor Oracle XA Transactions 15-20
Oracle XA Issues and Restrictions 15-21
Using Database Links in Oracle XA Applications 15-21
Managing Transaction Branches in Oracle XA Applications 15-22
Using Oracle XA with Oracle Real Application Clusters (Oracle RAC) 15-22
GLOBAL_TXN_PROCESSES Initialization Parameter 15-23
Managing Transaction Branches on Oracle RAC 15-23
Managing Instance Recovery in Oracle RAC with DTP Services (10.2) 15-24
Global Uniqueness of XIDs in Oracle RAC 15-25
Tight and Loose Coupling 15-25
SQL-Based Oracle XA Restrictions 15-26
Rollbacks and Commits 15-26
DDL Statements 15-26
Session State 15-26
EXEC SQL 15-26
Miscellaneous Restrictions 15-26

16 Developing Applications with the Publish-Subscribe Model

Introduction to the Publish-Subscribe Model 16-1
Publish-Subscribe Architecture 16-2
Database Events 16-2
Oracle Advanced Queuing 16-2
Client Notification 16-2
Publish-Subscribe Concepts 16-3
Examples of a Publish-Subscribe Mechanism 16-4

17 Using the Identity Code Package

Identity Concepts 17-1
What is the Identity Code Package? 17-5
Using the Identity Code Package 17-6
Storing RFID Tags in Oracle Database Using MGD_ID ADT 17-6
Creating a Table with MGD_ID Column Type and Storing EPC Tag Encodings in the Column 17-6
Constructing MGD_ID Objects to Represent RFID Tags 17-7
Inserting an MGD_ID Object into a Database Table 17-9
Querying MGD_ID Column Type 17-10
Building a Function-Based Index Using the Member Functions of the MGD_ID Column Type .. 17-10
Using MGD_ID ADT Functions 17-10

Using the get_component Function with the MGD_ID Object 17-11
Parsing Tag Data from Standard Representations 17-11
Reconstructing Tag Representations from Fields 17-12
Translating Between Tag Representations 17-13
Defining a Category of Identity Codes and Adding Encoding Schemes to an Existing Category 17-13
Creating a Category of Identity Codes 17-13
Adding Two Metadata Schemes to a Newly Created Category 17-14
Identity Code Package Types 17-18
DBMS_MGD_ID_UTL Package 17-18
Identity Code Metadata Tables and Views 17-19
Electronic Product Code (EPC) Concepts 17-21
RFID Technology and EPC v1.1 Coding Schemes 17-21
Product Code Concepts and Their Current Use 17-22
Electronic Product Code (EPC) 17-22
Global Trade Identification Number (GTIN) and Serializable Global Trade Identification Number (SGTIN) 17-24
Serial Shipping Container Code (SSCC) 17-24
Global Location Number (GLN) and Serializable Global Location Number (SGLN) . 17-24
Global Returnable Asset Identifier (GRAI) 17-24
Global Individual Asset Identifier (GIAI) 17-24
RFID EPC Network 17-24
Oracle Database Tag Data Translation Schema 17-24

18 Schema Object Dependency

Overview of Schema Object Dependencies 18-1
Querying Object Dependencies 18-4
Object Status 18-4
Invalidation of Dependent Objects 18-5
Session State and Referenced Packages 18-8
Security Authorization 18-8
Guidelines for Reducing Invalidation 18-8
Add Items to End of Package 18-8
Reference Each Table Through a View 18-9
Object Revalidation 18-9
Name Resolution in Schema Scope 18-10
Local Dependency Management 18-11
Remote Dependency Management 18-11
Dependencies Among Local and Remote Database Procedures 18-11
Dependencies Among Other Remote Objects 18-11
Dependencies of Applications 18-12
Remote Procedure Call (RPC) Dependency Management 18-12
Time-Stamp Dependency Mode 18-12
RPC-Signature Dependency Mode 18-13
Changing Names and Default Values of Parameters 18-15
Changing Specification of Parameter Mode IN 18-15
Changing Subprogram Body 18-15

Changing Data Type Classes of Parameters 18-16
Changing Package Types 18-17
Controlling Dependency Mode 18-18
Dependency Resolution 18-18
Suggestions for Managing Dependencies 18-19
Shared SQL Dependency Management 18-19

19 Edition-Based Redefinition

Editions 19-2
Editioned and Noneditioned Objects 19-2
Editionable and Noneditionable Schema Object Types 19-3
Rules for Editioned Objects 19-3
Enabling Editions for a User 19-4
Creating an Edition 19-5
Inherited and Actual Objects 19-5
Dropping Inherited Objects 19-7
Actualizing Referenced Objects 19-9
Making an Edition Available to Some Users 19-10
Making an Edition Available to All Users 19-10
Current Edition and Session Edition 19-10
Your Initial Session Edition 19-10
Changing Your Session Edition 19-11
Displaying the Names of the Current and Session Editions 19-12
When the Current Edition Might Differ from the Session Edition 19-12
Retiring an Edition 19-13
Dropping an Edition 19-14
Editioning Views 19-15
Creating an Editioning View 19-16
Partition-Extended Editioning View Names 19-16
Changing the 'Write-ability' of an Editioning View 19-17
Replacing an Editioning View 19-17
Dropping or Renaming the Base Table 19-17
Adding Indexes and Constraints to the Base Table 19-17
SQL Optimizer Index Hints 19-18
Crossedition Triggers 19-18
Forward Crossedition Triggers 19-18
Reverse Crossedition Triggers 19-19
Crossedition Trigger Interaction with Editions 19-19
Which Triggers Are Visible 19-19
What Kind of Triggers Can Fire 19-19
Firing Order 19-20
Crossedition Trigger Execution 19-21
Creating a Crossedition Trigger 19-22
Coding the Forward Crossedition Trigger Body 19-22
Transforming Data from Pre- to Post-Upgrade Representation 19-24
Preventing Lost Updates 19-26
Dropping the Crossedition Triggers 19-26

Displaying Information About Editions, Editioning Views, and Crossedition Triggers 19-27
Using Edition-Based Redefinition to Upgrade an Application 19-28
Preparing Your Application to Use Editioning Views 19-29
Procedure for Edition-Based Redefinition Using Only Editions 19-30
Procedure for Edition-Based Redefinition Using Editioning Views 19-32
Procedure for Edition-Based Redefinition Using Crossedition Triggers 19-33
Rolling Back the Application Upgrade 19-34
Reclaiming Space Occupied by Unused Table Columns 19-34
Example: Using Edition-Based Redefinition to Upgrade an Application 19-35
Existing Application 19-35
Preparing the Application to Use Editioning Views 19-37
Using Edition-Based Redefinition to Upgrade the Application 19-37

A Multithreaded extproc Agent

Why Use the Multithreaded extproc Agent? A-1
The Challenge of Dedicated Agent Architecture A-1
The Advantage of Multithreading A-1
Multithreaded extproc Agent Architecture A-2
Monitor Thread A-3
Dispatcher Threads A-4
Task Threads A-4
Administering the Multithreaded extproc Agent A-4
Agent Control Utility (agtctl) Commands A-5
Using agtctl in Single-Line Command Mode A-5
Setting Configuration Parameters for a Multithreaded extproc Agent A-6
Starting a Multithreaded extproc Agent A-6
Shutting Down a Multithreaded extproc Agent A-6
Examining the Value of Configuration Parameters A-7
Resetting a Configuration Parameter to Its Default Value A-7
Deleting an Entry for a Specific SID from the Control File A-7
Requesting Help A-7
Using Shell Mode Commands A-8
Example: Setting a Configuration Parameter A-8
Example: Starting a Multithreaded extproc Agent A-8
Configuration Parameters for Multithreaded extproc Agent Control A-8

Index

List of Examples

1–1 LOCK TABLE with SHARE MODE ... 1-13
1–2 How the Pro*COBOL Precompiler Uses Locks ... 1-23
1–3 Marking a Package Subprogram as Autonomous ... 1-36
1–4 AFTER SUSPEND Trigger Handles Suspended Storage Allocation ... 1-38
2–1 Displaying Current Date and Time with AD or BC Qualifier ... 2-11
2–2 Changing the Default Date Format ... 2-12
2–3 Changing the Default Time Format ... 2-12
2–4 Accessing Information in a SYS.ANYDATA Column ... 2-17
2–5 Querying the ROWID Pseudocolumn ... 2-22
3–1 Enforcing a Phone Number Format with Regular Expressions ... 3-10
3–2 Inserting Phone Numbers in Correct and Incorrect Formats ... 3-11
3–3 Using Back References to Reposition Characters ... 3-11
4–1 VENDOR_PARTS Table ... 4-4
4–2 Creating Indexes ... 4-6
4–3 Function-Based Index Allows Optimizer to Perform Range Scan ... 4-8
4–4 Function-Based Indexes ... 4-10
5–1 Inserting NULL Values into Columns with NOT NULL Constraints ... 5-3
5–2 Deferring Constraint Checks ... 5-10
5–3 Defining Constraints with the CREATE TABLE Statement ... 5-15
5–4 Defining Constraints with the ALTER TABLE Statement ... 5-15
5–5 Creating Enabled Constraints ... 5-17
5–6 Creating Disabled Constraints ... 5-18
5–7 Enabling Existing Constraints ... 5-18
5–8 Disabling Existing Constraints ... 5-19
5–9 Modifying Constraints ... 5-20
5–10 Renaming a Constraint ... 5-21
5–11 Dropping Constraints ... 5-22
5–12 Viewing Information About Constraints ... 5-24
6–1 Anonymous Block ... 6-2
6–2 Anonymous Block with Exception Handler for Predefined Error ... 6-3
6–3 Anonymous Block with Exception Handler for User-Defined Exception ... 6-3
6–4 Stored Procedure with Parameters ... 6-5
6–5 %TYPE and %ROWTYPE Attributes ... 6-7
6–6 Creating PL/SQL Package and Invoking Package Subprogram ... 6-11
6–7 Raising ORA-04068 ... 6-15
6–8 Trapping ORA-04068 ... 6-15
6–9 DML Statements that Reference Collections ... 6-16
6–10 SELECT Statements that Reference Collections ... 6-17
6–11 FOR Loops that Reference Collections and Return DML ... 6-18
6–12 Fetching Data with Cursor Variable ... 6-20
6–13 Cursor Variable with Discriminator ... 6-21
6–14 Compile-Time Errors ... 6-22
6–15 Invoking a Subprogram Interactively with SQL*Plus ... 6-30
6–16 Creating and Using a Session Variable with SQL*Plus ... 6-31
6–17 Invoking a Subprogram from Within Another Subprogram ... 6-32
6–18 PL/SQL Function in SQL Expression (Follows Rules) ... 6-37
6–19 PL/SQL Function in SQL Expression (Exception to Rule) ... 6-38
6–20 PRAGMA RESTRICT_REFERENCES ... 6-42
6–21 PRAGMA RESTRICT REFERENCES with TRUST on Invokee ... 6-43
6–22 PRAGMA RESTRICT REFERENCES with TRUST on Invoker ... 6-44
6–23 Overloaded Package Function with PRAGMA RESTRICT_REFERENCES ... 6-45
7–1 Is STANDARD and DBMS_STANDARD PL/Scope Identifier Data Available? ... 7-2
7–2 How Much Space is PL/Scope Data Using? ... 7-4
7–3 USAGE_CONTEXT_ID and USAGE_ID ... 7-6

7–4 Program Unit with Two Identifiers Named p 7-7
8–1 Profiling a PL/SQL Procedure 8-3
8–2 Invoking DBMS_HPROF.analyze 8-8
8–3 DBMSHP_RUNS Table for Sample PL/SQL Procedure 8-10
8–4 DBMSHP_FUNCTION_INFO Table for Sample PL/SQL Procedure 8-11
8–5 DBMSHP_PARENT_CHILD_INFO Table for Sample PL/SQL Procedure 8-11
8–6 Invoking DBMS_HPROF.analyze with Options 8-11
9–1 Creating and Configuring DADs 9-11
9–2 Authorizing DADs to be Created or Changed Later 9-12
9–3 Determining the Authentication Mode for a DAD 9-13
9–4 Showing the Authentication Mode for All DADs 9-14
9–5 Showing DAD Authorizations that Are Not in Effect 9-14
9–6 epgstat.sql Script Output for Example 9–1 9-15
9–7 Using HTP Functions to Generate HTML Tags 9-20
9–8 Using HTP.PRINT to Generate HTML Tags 9-20
9–9 HTML Drop-Down List 9-22
9–10 Passing Entry-Field Parameters from an HTML Form 9-23
9–11 Sending E-Mail from PL/SQL 9-26
9–12 Retrieving HTTP URL Contents from PL/SQL 9-27
10–1 simple.psp 10-1
10–2 Sample Returned HTML Page 10-6
10–3 simplewithuserinput.psp 10-8
10–4 Sample Comments in a PSP File 10-13
10–5 Loading PL/SQL Server Pages 10-13
10–6 Querying PL/SQL Server Page Source Code 10-14
10–7 show_prod_simple.psp 10-17
10–8 show_catalog_raw.psp 10-17
10–9 show_catalog_pretty.psp 10-18
10–10 show_product_partial.psp 10-18
10–11 show_product_highlighed.psp 10-19
10–12 product_form.psp 10-20
10–13 show_product_javascript.psp 10-20
11–1 Query to be Registered for Change Notification 11-2
11–2 Query Too Complex for QRCN in Guaranteed Mode 11-3
11–3 Query Whose Simplified Version Invalidates Objects 11-4
11–4 Creating a CQ_NOTIFICATION$_REG_INFO Object 11-21
11–5 Adding a Query to an Existing Registration 11-22
11–6 Creating Server-Side PL/SQL Notification Handler 11-28
11–7 Registering a Query 11-29
12–1 Retrieving a Lost Row with Oracle Flashback Query 12-7
12–2 Restoring a Lost Row After Oracle Flashback Query 12-7
12–3 Function that Can Return Row SCN from Table that has VPD 12-14
13–1 Pro*C/C++ Application 13-15
13–2 Pro*COBOL Application 13-17
15–1 xa_open String 15-8
15–2 Sample Open String Configuration 15-11
15–3 Transaction Started by an Application Server 15-14
15–4 Transaction Started by an Application Client 15-14
15–5 Using the DBMS_XA Package 15-16
18–1 Displaying Dependent and Referenced Object Types 18-1
18–2 Schema Object Change that Invalidates Some Dependents 18-3
18–3 View that Depends on Multiple Objects 18-4
18–4 Changing Body of Procedure get_hire_date 18-15
18–5 Changing Data Type Class of get_hire_date Parameter 18-17
18–6 Changing Names of Fields in Package Record Type 18-17

19–1 Inherited and Actual Objects 19-5
19–2 Dropping an Inherited Object 19-7
19–3 Creating an Object with the Name of a Dropped Inherited Object 19-8
19–4 Current Edition Differs from Session Edition 19-12
19–5 Crossedition Trigger that Handles Data Transformation Collisions 19-23
19–6 Edition-Based Redefinition of Very Simple Procedure 19-31
19–7 Creating the Existing Application 19-35
19–8 Viewing Data in Existing Table 19-36
19–9 Creating an Editioning View for the Existing Table 19-37
19–10 Creating Edition in Which to Upgrade the Application 19-37
19–11 Changing the Table and Replacing the Editioning View 19-38
19–12 Creating and Enabling the Crossedition Triggers 19-38
19–13 Applying the Transforms 19-41
19–14 Viewing Data in Changed Table 19-41
A–1 Setting Configuration Parameters and Starting agtctl A-4

List of Figures

1–1 Interaction Between Serializable Transaction and Another Transaction 1-26
1–2 Referential Integrity Check 1-27
1–3 Transaction Control Flow 1-31
1–4 Possible Sequences of Autonomous Transactions 1-32
1–5 Example: A Buy Order 1-33
1–6 Bank Withdrawal—Sufficient Funds 1-34
1–7 Bank Withdrawal—Insufficient Funds with Overdraft Protection 1-35
1–8 Bank Withdrawal—Insufficient Funds Without Overdraft Protection 1-36
5–1 Table with a UNIQUE Constraint 5-6
5–2 Tables with FOREIGN KEY Constraints 5-8
6–1 Exceptions and User-Defined Errors 6-25
9–1 PL/SQL Web Application 9-2
9–2 Processing Client Requests with Embedded PL/SQL Gateway 9-5
11–1 Middle-Tier Caching 11-8
11–2 Basic Process of Continuous Query Notification (CQN) 11-9
13–1 The OCI or OCCI Development Process 13-20
13–2 Software Layers 13-23
13–3 Objects and Their Relations 13-24
13–4 Supported Oracle Database Data Types 13-28
14–1 Oracle Database and External Procedures 14-27
15–1 Possible DTP Model 15-2
16–1 Oracle Publish-Subscribe Functionality 16-2
17–1 RFID Code Categories and Their Schemes 17-2
17–2 Oracle Database Tag Data Translation Markup Language Schema 17-4
A–1 Multithreaded extproc Agent Architecture A-3

List of Tables

1–1 COMMIT Statement Options 1-6
1–2 Use of COMMIT, SAVEPOINT, and ROLLBACK 1-8
1–3 Examples of Concurrency Under Explicit Locking 1-16
1–4 Ways to Display Locking Information 1-23
1–5 ANSI/ISO SQL Isolation Levels and Possible Transaction Interactions 1-25
1–6 ANSI/ISO SQL Isolation Levels Provided by Oracle Database 1-25
1–7 Comparison of READ COMMITTED and SERIALIZABLE Transactions 1-30
1–8 Possible Transaction Outcomes 1-33
2–1 SQL Character Data Types 2-2
2–2 Binary Floating-Point Format Components 2-6
2–3 Summary of Binary Format Storage Parameters 2-6
2–4 Range and Precision of Floating-Point Data Types 2-7
2–5 Special Values for Native Floating-Point Formats 2-7
2–6 Values Resulting from Exceptions 2-9
2–7 SQL Date and Time Data Types 2-10
2–8 Large Object Data Types 2-15
2–9 ANSI Data Type Conversions to Oracle Database Data Types 2-19
2–10 SQL/DS, DB2 Data Type Conversions to Oracle Database Data Types 2-20
2–11 Data Type Families 2-27
2–12 Display Types of SQL Built-In Functions 2-27
3–1 SQL Regular Expression Functions and Conditions 3-3
3–2 POSIX Metacharacters in Oracle Database Regular Expressions 3-5
3–3 POSIX and Multilingual Operator Relationships 3-8
3–4 PERL-Influenced Extensions in Oracle Database Regular Expressions 3-9
3–5 Pattern Matching Modifiers 3-10
3–6 Explanation of the Regular Expression Elements in Example 3–1 3-11
3–7 Explanation of the Regular Expression Elements in Example 3–3 3-12
6–1 Attributes of Subprogram Parameters 6-5
6–2 Parameter Modes 6-6
7–1 Identifier Types that PL/Scope Collects 7-8
7–2 Usages that PL/Scope Reports 7-9
8–1 Raw Profiler Output File Indicators 8-5
8–2 Function Names of Operations that the PL/SQL Hierarchical Profiler Tracks 8-6
8–3 PL/SQL Hierarchical Profiler Database Tables 8-7
8–4 DBMSHP_RUNS Table Columns 8-8
8–5 DBMSHP_FUNCTION_INFO Table Columns 8-9
8–6 DBMSHP_PARENT_CHILD_INFO Table Columns 8-10
9–1 Commonly Used Packages in the PL/SQL Web Toolkit 9-3
9–2 Mapping Between mod_plsql and Embedded PL/SQL Gateway DAD Attributes 9-7
9–3 Mapping Between mod_plsql and Embedded PL/SQL Gateway Global Attributes 9-8
9–4 Authentication Possibilities for a DAD 9-11
10–1 PSP Elements 10-4
11–1 Continuous Query Notification Registration Options 11-11
11–2 Attributes of CQ_NOTIFICATION$_REG_INFO 11-19
11–3 Quality-of-Service Flags 11-21
11–4 Attributes of CQ_NOTIFICATION$_DESCRIPTOR 11-25
11–5 Attributes of CQ_NOTIFICATION$_TABLE 11-26
11–6 Attributes of CQ_NOTIFICATION$_QUERY 11-26
11–7 Attributes of CQ_NOTIFICATION$_ROW 11-27
12–1 Oracle Flashback Version Query Row Data Pseudocolumns 12-8
12–2 Flashback TRANSACTION_BACKOUT Options 12-17
12–3 Static Data Dictionary Views for Flashback Data Archive Files 12-23
13–1 PL/SQL Packages and Their Java Equivalents 13-13

14–1 Parameter Data Type Mappings 14-16
14–2 External Data Type Mappings 14-17
14–3 Properties and Data Types 14-21
15–1 Required XA Features Published by Oracle Database 15-5
15–2 XA Library Subprograms 15-5
15–3 Oracle XA Interface Extensions 15-6
15–4 Required Fields of xa_open string 15-9
15–5 Optional Fields in the xa_open String 15-9
15–6 TX Interface Functions 15-13
15–7 TPM Replacement Statements 15-15
15–8 Sample Trace File Contents 15-19
15–9 Tightly and Loosely Coupled Transaction Branches 15-22
17–1 General Structure of EPC Encodings 17-2
17–2 Identity Code Package ADTs 17-18
17–3 MGD_ID ADT Subprograms 17-18
17–4 DBMS_MGD_ID_UTL Package Utility Subprograms 17-19
17–5 Definition and Description of the MGD_ID_CATEGORY Metadata View 17-20
17–6 Definition and Description of the USER_MGD_ID_CATEGORY Metadata View 17-20
17–7 Definition and Description of the MGD_ID_SCHEME Metadata View 17-21
17–8 Definition and Description of the USER_MGD_ID_SCHEME Metadata View 17-21
18–1 Database Object Status 18-4
18–2 Operations that Cause Fine-Grained Invalidation 18-5
18–3 Data Type Classes 18-16
19–1 *_ Dictionary Views with Edition Information 19-27
19–2 *_ Dictionary Views with Editioning View Information 19-28
A–1 Agent Control Utility (agtctl) Commands A-5
A–2 Configuration Parameters for agtctl A-8

Preface

Oracle Database Advanced Application Developer's Guide explains topics that experienced application developers reference repeatedly. Information in this guide applies to features that work the same on all supported platforms, and does not include system-specific information.

Preface topics:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

Oracle Database Advanced Application Developer's Guide is intended for application developers who are either developing applications or converting applications to run in the Oracle Database environment. This guide is also valuable to anyone who is interested in the development of database applications, such as systems analysts and project managers.

To use this document effectively, you need a working knowledge of:

- Application programming
- Structured Query Language (SQL)
- Object-oriented programming

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at `http://www.oracle.com/accessibility/`.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/support/contact.html` or visit `http://www.oracle.com/accessibility/support.html` if you are hearing impaired.

Related Documents

For more information, see these documents in the Oracle Database 11*g* Release 2 (11.2) documentation set:

- *Oracle Database PL/SQL Language Reference*
- *Oracle Call Interface Programmer's Guide*
- *Oracle Database Security Guide*
- *Pro*C/C++ Programmer's Guide*
- *Oracle Database SQL Language Reference*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle XML Developer's Kit Programmer's Guide*
- *Oracle XML DB Developer's Guide*
- *Oracle Database Globalization Support Guide*
- *Oracle Database Sample Schemas*

See also:

- *Oracle PL/SQL Tips and Techniques* by Joseph C. Trezzo. Oracle Press, 1999.
- *Oracle PL/SQL Programming* by Steven Feuerstein. 3rd Edition. O'Reilly & Associates, 2002.
- *Oracle PL/SQL Developer's Workbook* by Steven Feuerstein. O'Reilly & Associates, 2000.
- *Oracle PL/SQL Best Practices* by Steven Feuerstein. O'Reilly & Associates, 2001.

Conventions

This document uses these text conventions:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Also:

- `*_view` means all static data dictionary views whose names end with `view`. For example, `*_ERRORS` means `ALL_ERRORS`, `DBA_ERRORS`, and `USER_ERRORS`. For more information about any static data dictionary view, or about static dictionary views in general, see *Oracle Database Reference*.
- Table names not qualified with schema names are in the sample schema `HR`. For information about the sample schemas, see *Oracle Database Sample Schemas*.

What's New in Application Development?

This topic briefly describes the new Oracle Database features that this book documents and provides links to more information.

Topics:

- Oracle Database 11g Release 2 (11.2.0.2) Feature
- Oracle Database 11g Release 2 (11.2.0.1) Features
- Oracle Database 11g Release 1 (11.1) Features

Oracle Database 11*g* Release 2 (11.2.0.2) Feature

Edition Attribute of Database Service

Before Release 11.2.0.2, you could not specify your initial session edition when using a database service to connect to Oracle Database. If you wanted to use **Edition-Based Redefinition** for hot rollover, where some database clients use the pre-upgrade edition while others use the post-upgrade edition, then you had to change the client code.

As of Release 11.2.0.2, you can specify the initial session edition as an attribute of a database service, which makes it easier to ensure that each session uses the desired edition during hot rollover. For more information, see "Your Initial Session Edition" on page 19-10.

As of Release 11.2.0.2, each `*_SERVICES` static data dictionary view has an `EDITION` column that shows the default initial session edition. For more information, see "Displaying Information About Editions, Editioning Views, and Crossedition Triggers" on page 19-27.

Oracle Database 11*g* Release 2 (11.2.0.1) Features

The Oracle Database features for Release 11.2.0.1 are:

- Flashback Transaction Foreign Key Dependency Tracking
- Fine-Grained Invalidation for Triggers
- Edition-Based Redefinition
- APPLYING_CROSSEDITION_TRIGGER Function
- IGNORE_ROW_ON_DUPKEY_INDEX Hint
- CHANGE_DUPKEY_ERROR_INDEX Hint
- DBMS_PARALLEL_EXECUTE Package

- Internet Protocol version 6 (IPv6) Support

Flashback Transaction Foreign Key Dependency Tracking

Flashback Transaction (the `DBMS_FLASHBACK.TRANSACTION_BACKOUT` procedure) with the `CASCADE` option rolls back a transaction and its dependent transactions while the database remains online.

Before Release 11.2.0.1, Flashback Transaction did not track foreign key dependencies. Therefore, if you tried to use Flashback Transaction with the `CASCADE` option to roll back a transaction that had foreign key dependencies, you could get a foreign key violation error. The workaround was to include the foreign-key-dependent transactions in the list of transactions to roll back.

As of Release 11.2.0.1, when using Flashback Transaction with the `CASCADE` option, you do not have to include any dependent transactions in the list of transactions to be rolled back.

Foreign key dependency tracking for Flashback Transaction requires that you enable foreign key supplemental logging. For instructions, see "Configuring Your Database for Flashback Transaction" on page 12-4. For information about Flashback Transaction, see "Using Flashback Transaction" on page 12-15.

Fine-Grained Invalidation for Triggers

The Release 11.1 feature "Fine-Grained Invalidation" on page xxxix has been extended to triggers.

Edition-Based Redefinition

Edition-based redefinition enables you to upgrade the database component of an application while it is in use, thereby minimizing or eliminating down time.

To upgrade an application while it is in use, you copy the database objects that comprise the application and redefine the copied objects in isolation. Your changes do not affect users of the application—they continue to run the unchanged application. When you are sure that your changes are correct, you make the upgraded application available to all users.

Using edition-based redefinition means using one or more of its component features. The features you use, and the down time, depend on these factors:

- What kind of database objects you redefine
- How available the database objects must be to users while you are redefining them
- Whether you make the upgraded application available to some users while others continue to use the older version of the application

You always use the **edition** feature to copy the database objects and redefine the copied objects in isolation.

If you change the structure of one or more tables, you also use the feature **editioning views**.

If other users must be able to change data in the tables while you are changing their structure, you also use **forward crossedition triggers**. If the pre- and post-upgrade applications will be in ordinary use at the same time (hot rollover), you also use **reverse crossedition triggers**. Crossedition triggers are not a permanent part of the application—you drop them when all users are using the post-upgrade application.

For more information, see Chapter 19, "Edition-Based Redefinition."

APPLYING_CROSSEDITION_TRIGGER Function

The body of a forward crossedition trigger must handle data transformation collisions. If your collision-handling strategy depends on why the trigger is running, you can determine the reason with the function `APPLYING_CROSSEDITION_TRIGGER`, which is defined in the package `DBMS_STANDARD`.

For more information, see "Handling Data Transformation Collisions" on page 19-23.

IGNORE_ROW_ON_DUPKEY_INDEX Hint

When a statement of the form `INSERT INTO` *`target subquery`* runs, a unique key for some rows to be inserted might collide with existing rows. Suppose that your application must ignore such collisions and insert the rows that do not collide with existing rows.

Before Release 11.2.0.1, you had to write a PL/SQL program which, in a block with a `NULL` handler for the `DUP_VAL_ON_INDEX` exception, selected the source rows and then inserted them, one at a time, into the target.

As of Release 11.2.0.1, you do not have to write a PL/SQL program. You can use the `IGNORE_ROW_ON_DUPKEY_INDEX` hint in an `INSERT` statement, which is easier to write and runs much faster. This hint is especially helpful when implementing crossedition triggers.

For more information, see "Handling Data Transformation Collisions" on page 19-23.

CHANGE_DUPKEY_ERROR_INDEX Hint

When an `INSERT` or `UPDATE` statement runs, a unique key might collide with existing rows.

Before Release 11.2.0.1, the collision caused error ORA-00001. You could tell that a collision had occurred, but you could not tell where.

As of Release 11.2.0.1, you can use the `CHANGE_DUPKEY_ERROR_INDEX` hint in an `INSERT` or `UPDATE` statement, specifying that when a unique key violation occurs for a specified index or set of columns, ORA-38911 is reported instead of ORA-00001. This hint is especially helpful when implementing crossedition triggers.

For more information, see "Handling Data Transformation Collisions" on page 19-23.

DBMS_PARALLEL_EXECUTE Package

The `DBMS_PARALLEL_EXECUTE` package enables you to incrementally update the data in a large table in parallel, in two high-level steps:

1. Group sets of rows in the table into smaller chunks.
2. Apply the desired `UPDATE` statement to the chunks in parallel, committing each time you have finished processing a chunk.

This technique improves performance, reduces rollback space consumption, and reduces the number of row locks held. The `DBMS_PARALLEL_EXECUTE` package is recommended whenever you are updating a lot of data; for example, when you are applying forward crossedition triggers.

For more information, see "Transforming Data from Pre- to Post-Upgrade Representation" on page 19-24.

Internet Protocol version 6 (IPv6) Support

Internet Protocol version 6 (IPv6) supports a much larger address space than IPv4 does. An IPv6 address has 128 bits, while an IPv4 address has only 32 bits.

Applications that use network addresses might need small changes, and recompilation, to accommodate IPv6 addresses. For more information, see "Performing Network Operations in PL/SQL Subprograms" on page 9-25.

The agent control utility, `agtctl`, which starts a multithreaded `extproc` agent, now accepts IPv6 addresses. For more information, see "Configuration Parameters for Multithreaded extproc Agent Control" on page A-8.

> **See Also:** *Oracle Database Net Services Administrator's Guide* for detailed information about IPv6 support in Oracle Database

Oracle Database 11*g* Release 1 (11.1) Features

The application development features for Release 11.1 are:

- WAIT Option for Data Definition Language (DDL) Statements
- Binary XML Support for Oracle XML Database
- Metadata for SQL Built-In Functions
- Enhancements to Regular Expression Built-in Functions
- Invisible Indexes
- PL/SQL Function Result Cache
- Sequences in PL/SQL Expressions
- PL/Scope
- PL/SQL Hierarchical Profiler
- Query Result Change Notification
- Flashback Transaction
- Flashback Data Archive (Oracle Total Recall)
- XA API Available Within PL/SQL
- Support for XA/JTA in Oracle Real Application Clusters (Oracle RAC) Environment
- Identity Code Package
- Enhanced Online Index Creation and Rebuilding
- Embedded PL/SQL Gateway
- Oracle Database Spawns Multithreaded extproc Agent Directly by Default
- Fine-Grained Invalidation

WAIT Option for Data Definition Language (DDL) Statements

DDL statements require exclusive locks on internal structures. If these locks are unavailable when a DDL statement is issued, the DDL statement fails, though it might have succeeded if it had been issued subseconds later. The `WAIT` option of the SQL statement `LOCK TABLE` enables a DDL statement to wait for its locks for a specified period before failing.

For more information, see "Choosing a Locking Strategy" on page 1-12.

Binary XML Support for Oracle XML Database

Binary XML is a third way to represent an XML document. Binary XML complements, rather than replaces, the existing object-relational storage and `CLOB` storage representations. Binary XML has two significant benefits:

- XML operations can be significantly optimized, with or without an XML schema is available.
- The internal representation of XML is the same on disk, in memory, and on wire.

As with other storage mechanisms, the details of binary XML storage are transparent to you. You continue to use `XMLType` and its associated methods and operators.

For more information, see "Representing XML" on page 2-16.

See Also: *Oracle XML DB Developer's Guide*

Metadata for SQL Built-In Functions

Metadata for SQL built-in functions is accessible through dynamic performance (`V$`) views. Third-party tools can leverage built-in SQL functions without maintaining their metadata in the application layer.

For more information, see "Metadata for SQL Built-In Functions" on page 2-26.

Enhancements to Regular Expression Built-in Functions

The regular expression built-in functions `REGEXP_INSTR` and `REGEXP_SUBSTR` have increased functionality. A new regular expression built-in function, `REGEXP_COUNT`, returns the number of times a pattern appears in a string. These functions act the same in SQL and PL/SQL.

For more information, see "Oracle Database Implementation of Regular Expressions" on page 3-2.

See Also: *Oracle Database SQL Language Reference*

Invisible Indexes

An invisible index is maintained by Oracle Database for every data manipulation language (DML) statement, but is ignored by the optimizer unless you explicitly set the parameter `OPTIMIZER_USE_INVISIBLE_INDEXES` to `TRUE` on a session or system level.

Making an index invisible is an alternative to making it unusable or dropping it. Using invisible indexes, you can:

- Test the removal of an index before dropping it
- Create invisible indexes temporarily for specialized, nonstandard operations, such as online application upgrades, without affecting the behavior of existing applications

For more information, see "Drop Unused Indexes" on page 4-5.

PL/SQL Function Result Cache

Before Release 11.1, if you wanted your PL/SQL application to cache the results of a function, you had to design and code the cache and cache-management subprograms. If multiple sessions ran your application, each session had to have its own copy of the cache and cache-management subprograms. Sometimes each session had to perform the same expensive computations.

As of Release 11.1, PL/SQL provides a function result cache. Because the function result cache is stored in a shared global area (SGA), it is available to any session that runs your application.

For more information, see "PL/SQL Function Result Cache" on page 6-10.

> **See Also:** *Oracle Database PL/SQL Language Reference*

Sequences in PL/SQL Expressions

The pseudocolumns `CURRVAL` and `NEXTVAL` make writing PL/SQL source code easier for you and improve run-time performance and scalability. You can use *`sequence_name`*`.CURRVAL` and *`sequence_name`*`.NEXTVAL` wherever you can use a `NUMBER` expression.

See Example 6–6 on page 6-11.

> **See Also:** *Oracle Database PL/SQL Language Reference*

PL/Scope

PL/Scope is a compiler-driven tool that collects and organizes data about user-defined identifiers from PL/SQL source code. Because PL/Scope is a compiler-driven tool, you use it through interactive development environments (such as SQL Developer and JDeveloper), rather than directly.

PL/Scope enables the development of powerful and effective PL/Scope source code browsers that increase PL/SQL developer productivity by minimizing time spent browsing and understanding source code.

For a detailed description of PL/Scope, see Chapter 7, "Using PL/Scope."

PL/SQL Hierarchical Profiler

Nonhierarchical (**flat**) profilers record the time that a program spends within each subprogram—the **function time** or **self time** of each subprogram. Function time is helpful, but often inadequate. For example, it is helpful to know that a program spends 40% of its time in the subprogram `INSERT_ORDER`, but it is more helpful to know which subprograms call `INSERT_ORDER` often and the total time the program spends under `INSERT_ORDER` (including its descendent subprograms). Hierarchical profilers provide such information.

The PL/SQL hierarchical profiler does this:

- Reports the dynamic execution profile of your PL/SQL program, organized by subprogram calls
- Accounts for SQL and PL/SQL execution times separately
- Requires no special source or compile-time preparation
- Stores results in database tables (**hierarchical profiler tables**) for custom report generation by integrated development environment (IDE) tools (such as SQL Developer and third-party tools)

  To generate simple HTML reports from raw profiler output, you can use the `plshprof` command-line utility.

Each subprogram-level summary in the dynamic execution profile includes information such as:

- Number of calls to the subprogram
- Time spent in the subprogram itself (**function time** or **self time**)

- Time spent in the subprogram itself and in its descendent subprograms (**subtree time**)
- Detailed parent-children information, for example:
  - All callers of a given subprogram (parents)
  - All subprograms that a given subprogram called (children)
  - How much time was spent in subprogram *x* when called from *y*
  - How many calls to subprogram *x* were from *y*

You can browse the generated HTML reports in any browser. The browser's navigational capabilities, combined with well chosen links, provide a powerful way to analyze performance of large applications, improve application performance, and lower development costs.

For a detailed description of PL/SQL hierarchical profiler, see Chapter 8, "Using the PL/SQL Hierarchical Profiler."

Query Result Change Notification

Before Release 11.1, Continuous Query Notification (CQN) published only object change notifications, which result from DML or DDL changes to the objects associated with registered the queries.

As of Release 11.1, CQN can also publish query result change notifications, which result from DML or DDL changes to the result set associated with the registered queries. New static data dictionary views enable you to see which queries are registered for result-set-change notifications (see "Querying CQN Registrations" on page 11-24).

For more information, see Chapter 11, "Using Continuous Query Notification (CQN)."

Flashback Transaction

The `DBMS_FLASHBACK.TRANSACTION_BACKOUT` procedure rolls back a transaction and its dependent transactions while the database remains online. This recovery operation uses undo data to create and run the compensating transactions that return the affected data to its original state.

For more information, see "Using Flashback Transaction" on page 12-15.

Flashback Data Archive (Oracle Total Recall)

A Flashback Data Archive provides the ability to store and track transactional changes to a record over its lifetime. It is no longer necessary to build this intelligence into the application. A Flashback Data Archive is useful for compliance with record stage policies and audit reports.

For more information, see "Using Flashback Data Archive (Oracle Total Recall)" on page 12-18.

XA API Available Within PL/SQL

The XA interface functionality that supports transactions involving multiple resource managers, such as databases and queues, is now available within PL/SQL. You can use PL/SQL to switch and share transactions across SQL*Plus sessions and across processes.

For more information, see "Using the DBMS_XA Package" on page 15-16.

Support for XA/JTA in Oracle Real Application Clusters (Oracle RAC) Environment

An XA transaction now spans Oracle RAC instances by default, enabling any application that uses XA to take full advantage of the Oracle RAC environment, enhancing the availability and scalability of the application.

For more information, see "Using Oracle XA with Oracle Real Application Clusters (Oracle RAC)" on page 15-22.

Identity Code Package

The Identity Code Package provides tools to store, retrieve, encode, decode, and translate between various product or identity codes, including Electronic Product Code (EPC), in Oracle Database. The Identity Code Package provides new data types, metadata tables and views, and PL/SQL packages for storing EPC standard RFID tags or new types of RFID tags in a user table.

The Identity Code Package enables Oracle Database to recognize EPC coding schemes, to support efficient storage and component-level retrieval of EPC data, and to meet the EPCglobal Tag Data Translation 1.0 (TDT) standard that defines how to decode, encode, and translate between various EPC RFID tag representations.

The Identity Code Package also provides an extensible framework that enables you to use pre-existing coding schemes with applications that are not included in the EPC standard and adapt Oracle Database both to these older systems and to evolving identity codes that might become part of a future EPC standard.

The Identity Code Package also lets you create your own identity codes by first registering the new encoding category, registering the new encoding type, and then registering the new components associated with each new encoding type.

For more information, see Chapter 17, "Using the Identity Code Package."

Enhanced Online Index Creation and Rebuilding

Online index creation and rebuilding no longer requires a DML-blocking lock.

Before Release 11.1, online index creation and rebuilding required a very short-term DML-blocking lock at the end of the rebuilding. The DML-blocking lock could cause a spike in the number of waiting DML operations, and therefore a short drop and spike of system usage. This system usage anomaly could trigger operating system alarm levels.

Embedded PL/SQL Gateway

The PL/SQL gateway enables a user-written PL/SQL subprogram to be invoked in response to a URL with parameters derived from an HTTP request. `mod_plsql` is a form of the gateway that exists as a plug-in to the Oracle HTTP Server. Now the PL/SQL gateway is also embedded in the database itself. The embedded PL/SQL gateway uses the internal Oracle XML Database Listener and does not depend on the Oracle HTTP Server. You configure the embedded version of the gateway with the `DBMS_EPG` package.

For more information, see "Using Embedded PL/SQL Gateway" on page 9-4.

Oracle Database Spawns Multithreaded extproc Agent Directly by Default

When an application calls an external C procedure, either Oracle Database or Oracle Listener starts the external procedure agent, `extproc`.

Before Release 11.1, Oracle Listener spawned the multithreaded `extproc` agent, and you defined environment variables for `extproc` in the file `listener.ora`.

As of Release 11.1, by default, Oracle Database spawns `extproc` directly, eliminating the risk that Oracle Listener might spawn `extproc` unexpectedly. This default configuration is recommended for maximum security. If you use it, you define environment variables for `extproc` in the file `extproc.ora`.

For more information, including situations in which you cannot use the default configuration, see "Loading External Procedures" on page 14-4.

Fine-Grained Invalidation

Before Release 11.1, a DDL statement that changed a referenced object invalidated all of its dependents.

As of Release 11.1, a DDL statement that changes a referenced object invalidates only the dependents for which either of these statements is true:

- The dependent relies on the attribute of the referenced object that the DDL statement changed.
- The compiled metadata of the dependent is no longer correct for the changed referenced object.

For example, if view `v` selects columns `c1` and `c2` from table `t`, a DDL statement that changes only column `c3` of `t` does not invalidate `v`.

For more information, see "Invalidation of Dependent Objects" on page 18-5.

Part I

SQL for Application Developers

This part presents information that application developers need about Structured Query Language (SQL), which is used to manage information in an Oracle Database.

Chapters:

- Chapter 1, "SQL Processing for Application Developers"
- Chapter 2, "Using SQL Data Types in Database Applications"
- Chapter 3, "Using Regular Expressions in Database Applications"
- Chapter 4, "Using Indexes in Database Applications"
- Chapter 5, "Maintaining Data Integrity in Database Applications"

> **See Also:** *Oracle Database SQL Language Reference* for a complete description of SQL

1

SQL Processing for Application Developers

This chapter explains what application developers must know about how Oracle Database processes SQL statements. Before reading this chapter, read the basic information about SQL processing in *Oracle Database Concepts*.

Topics:

- Description of SQL Statement Processing
- Grouping Operations into Transactions
- Ensuring Repeatable Reads with Read-Only Transactions
- Using Cursors
- Locking Tables Explicitly
- Using Oracle Lock Management Services (User Locks)
- Using Serializable Transactions for Concurrency Control
- Autonomous Transactions
- Resuming Execution After Storage Allocation Errors

Description of SQL Statement Processing

This topic explains what happens during each stage of processing the execution of a SQL statement, using a DML statement as an example.

Assume that you are using a Pro*C program to increase the salary for all employees in a department. The program you are using has connected to Oracle Database and you are connected to the proper schema to update the `employees` table. You can embed this SQL statement in your program:

```
EXEC SQL UPDATE employees SET salary = 1.10 * salary
  WHERE department_id = :department_id;
```

`Department_id` is a program variable containing a value for department number. When the SQL statement runs, it uses the value of `department_id` that the application program provides.

Stages of SQL Statement Processing

> **Note:** DML statements use all stages. Transaction management, session management, and system management SQL statements use only stages 2 and 8.

1. **Open or create a cursor.**

   A program interface call opens or creates a cursor, in expectation of a SQL statement. Most applications create the cursor implicitly (automatically). Precompiler programs can create the cursor either implicitly or explicitly.

2. **Parse the statement.**

   The user process passes the SQL statement to Oracle Database, which loads a parsed representation of the statement into the shared SQL area. Many errors can be caught during this stage of statement processing.

   > **Note:** For a DDL statement, parsing includes data dictionary lookup and execution.

   See Also:

   - *Oracle Database Concepts* for information about parsing
   - "Shared SQL Areas" on page 1-3

3. **Determine if the statement is a query.**

4. **If the statement is a query, describe its results.**

   > **Note:** This stage is necessary only if the characteristics of the result are unknown; for example, when a user enters the query interactively.

   Oracle Database determines the characteristics (data types, lengths, and names) of the result.

5. **If the statement is a query, define its output.**

   You specify the location, size, and data type of variables defined to receive each fetched value. These variables are called **define variables**. Oracle Database performs data type conversion if necessary.

   See Also: *Oracle Database Concepts* for information about the `DEFINE` stage

6. **Bind any variables.**

   Oracle Database knows the meaning of the SQL statement but does not have enough information to run it. Oracle Database needs values for any variables in the statement. In the example, Oracle Database needs a value for `department_id`. The process of obtaining these values is called **binding variables**.

A program must specify the location (memory address) of the value. End users of applications may be unaware that they are specifying bind variables, because the Oracle Database utility can prompt them for values.

Because you specify the location (binding by reference), you need not rebind the variable before rerunning the statement, even if you change its value. Each time Oracle Database runs the statement, it gets the value of the variable from its address.

You must also specify a data type and length for each value (unless they are implied or defaulted) if Oracle Database must perform data type conversion.

> **See Also:** For more information about specifying a data type and length for a value:
>
> - *Oracle Call Interface Programmer's Guide*
> - *Pro*C/C++ Programmer's Guide*

7. **(Optional) Parallelize the statement.**

   Oracle Database can parallelize queries and some data definition language (DDL) operations such as index creation, creating a table with a subquery, and operations on partitions. Parallelization causes multiple server processes to perform the work of the SQL statement so it can complete faster.

8. **Run the statement.**

   Oracle Database runs the statement. If the statement is a query or an `INSERT` statement, the database locks no rows, because no data is changing. If the statement is an `UPDATE` or `DELETE` statement, the database locks all rows that the statement affects, until the next `COMMIT`, `ROLLBACK`, or `SAVEPOINT` for the transaction, thereby ensuring data integrity.

   For some statements, you can specify multiple executions to be performed. This is called **array processing**. Given *n* number of executions, the bind and define locations are assumed to be the beginning of an array of size *n*.

9. **If the statement is a query, fetch its rows.**

   Oracle Database selects rows and, if the query has an `ORDER BY` clause, orders the rows. Each successive fetch retrieves another row of the result set, until the last row has been fetched.

10. **Close the cursor.**

    Oracle Database closes the cursor.

> **Note:** To rerun a transaction management, session management, or system management SQL statement, use another `EXECUTE` statement.

Shared SQL Areas

Oracle Database automatically notices when applications send similar SQL statements to the database. The SQL area used to process the first occurrence of the statement is *shared*—that is, used for processing subsequent occurrences of that same statement. Therefore, only one shared SQL area exists for a unique statement. Because shared SQL areas are shared memory areas, any Oracle Database process can use a shared SQL area. The sharing of SQL areas reduces memory use on the database server, thereby increasing system throughput.

In determining whether statements are similar or identical, Oracle Database considers both SQL statements issued directly by users and applications and recursive SQL statements issued internally by DDL statements.

> **See Also:** *Oracle Database Performance Tuning Guide* for more information about shared SQL

Grouping Operations into Transactions

Topics:

- Deciding How to Group Operations in Transactions
- Improving Transaction Performance
- Committing Transactions
- Managing Commit Redo Action
- Rolling Back Transactions
- Defining Transaction Savepoints

Deciding How to Group Operations in Transactions

Typically, deciding how to group operations in transactions is the concern of application designers who use programming interfaces to Oracle Database. When deciding how to group transactions:

- Define transactions such that work is accomplished in logical units and data remains consistent.
- Ensure that data in all referenced tables is in a consistent state before the transaction begins and after it ends.
- Ensure that each transaction consists only of the SQL statements or PL/SQL blocks that comprise one consistent change to the data.

For example, suppose that you write a Web application that enables users to transfer funds between accounts. The transaction must include the debit to one account, executed by one SQL statement, and the credit to another account, executed by another SQL statement. Both statements must fail or succeed as a unit of work; one statement must not be committed without the other. Do not include unrelated actions, such as a deposit to one account, in the transaction.

Improving Transaction Performance

As an application developer, you must try to improve performance. Consider using these performance enhancement techniques when designing and writing your application:

- If you use rollback segments for undo, then use the `SET TRANSACTION` statement with the `USE ROLLBACK SEGMENT` clause to explicitly assign a transaction to a rollback segment.

  This technique can eliminate the need to allocate additional extents dynamically. Allocating extents dynamically can reduce system performance.

> **Note:** If you use automatic undo management, then Oracle Database ignores the `USE ROLLBACK SEGMENT` clause.

- Establish standards for writing SQL statements so that you can take advantage of shared SQL areas.

  Oracle Database recognizes identical SQL statements and enables them to share memory areas, reducing memory usage on the database server and increasing system throughput.

- Collect statistics that Oracle Database can use to implement a cost-based approach to SQL statement optimization, and use additional hints to the optimizer as needed.

  To collect most statistics, use the `DBMS_STATS` package, which lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. For more information about this package, see *Oracle Database PL/SQL Packages and Types Reference*.

  To collect statistics unrelated to the cost-based optimizer (such as information about free list blocks), use the SQL statement `ANALYZE`. For more information about this statement, see *Oracle Database SQL Language Reference*.

  For more information about hints, see *Oracle Database SQL Language Reference*.

- Before beginning a transaction, invoke `DBMS_APPLICATION_INFO` procedures to record the name of the transaction in the database for later use when tracking its performance with Oracle Trace and the SQL trace facility. For information about the `DBMS_APPLICATION_INFO` package, see *Oracle Database PL/SQL Packages and Types Reference*.

- Increase user productivity and query efficiency by including user-written PL/SQL functions in SQL expressions. For details, see "Invoking Stored PL/SQL Functions from SQL Statements" on page 6-35.

- When writing a PL/SQL application, create explicit cursors. For information, see "Using Cursors" on page 1-9.

- Reduce parsing frequency and improve performance in precompiler programs by increasing the number of cursors with `MAX_OPEN_CURSORS`.

- Use the `SET TRANSACTION` statement with the `ISOLATION LEVEL` set to `SERIALIZABLE` to get ANSI/ISO serializable transactions.

> **See Also:** *Oracle Database Concepts* for more information about transaction management

Committing Transactions

To commit a transaction, use the `COMMIT` statement. These two statements are equivalent and commit the current transaction:

```
COMMIT WORK;
COMMIT;
```

The `COMMIT` statements lets you include the `COMMENT` parameter along with a comment that provides information about the transaction being committed. This option is useful for including information about the origin of the transaction when you commit distributed transactions:

```
COMMIT COMMENT 'Dallas/Accts_pay/Trans_type 10B';
```

See Also: *Oracle Database SQL Language Reference* for information about the COMMIT statement

Managing Commit Redo Action

When a transaction updates Oracle Database, it generates a corresponding redo entry. Oracle Database buffers the redo entry to the redo log until the transaction completes. When the transaction commits, the log writer process (LGWR) writes redo records to disk for the buffered redo entries of all changes in the transaction. By default, Oracle Database writes the redo entries to disk before the call returns to the client. This action introduces a latency in the commit, because the application must wait for the redo entries to be persistent on disk.

Oracle Database lets you change the handling of commit redo to fit the needs of your application. If your application requires very high transaction throughput and you are willing to trade commit durability for lower commit latency, then you can change the default COMMIT options so that the application need not wait for the database to write data to the online redo logs.

Table 1–1 describes the COMMIT options.

Caution: With the NOWAIT option, a failure that occurs after the commit message is received, but before the redo log records are written, can falsely indicate to a transaction that its changes are persistent.

Table 1–1 COMMIT Statement Options

Option	Effect
WAIT (default)	Ensures that the COMMIT statement returns only after the corresponding redo information is persistent in the online redo log. When the client receives a successful return from this COMMIT statement, the transaction has been committed to durable media. A failure that occurs after a successful write to the log might prevent the success message from returning to the client, in which case the client cannot tell whether the transaction committed.
NOWAIT (alternative to WAIT)	The COMMIT statement returns to the client regardless of whether the write to the redo log has completed. This behavior can increase transaction throughput.
BATCH (alternative to IMMEDIATE)	Buffers the redo information to the redo log, along with other concurrently running transactions. After collecting sufficient redo information, initiates a disk write to the redo log. This behavior is called **group commit**, because it writes redo information for multiple transactions to the log in a single I/O operation.
IMMEDIATE (default)	LGWR writes the transaction redo information to the log. Because this operation option forces a disk I/O, it can reduce transaction throughput.

To change the COMMIT options, use either the COMMIT statement (described in *Oracle Database SQL Language Reference*) or the appropriate initialization parameter. For information about initialization parameters, see *Oracle Database Reference*.

Note: You cannot change the default IMMEDIATE and WAIT action for distributed transactions.

If your application uses Oracle Call Interface (OCI), then you can modify redo action by setting these flags in the OCITransCommit function within your application:

- OCI_TRANS_WRITEWAIT
- OCI_TRANS_WRITENOWAIT
- OCI_TRANS_WRITEBATCH
- OCI_TRANS_WRITEIMMED

Caution: OCI_TRANS_WRITENOWAIT can cause silent transaction loss with shutdown termination, startup force, and any instance or node failure. On an Oracle RAC system, asynchronously committed changes might not be immediately available to read on other instances.

See Also: *Oracle Call Interface Programmer's Guide* for information about the OCITransCommit function

The specification of the NOWAIT and BATCH options has a small window of vulnerability in which Oracle Database can roll back a transaction that your application views as committed. Your application must be able to tolerate these scenarios:

- The database host fails, which causes the database to lose redo entries that were buffered but not yet written to the online redo logs.
- A file I/O problem prevents LGWR from writing buffered redo entries to disk. If the redo logs are not multiplexed, then the commit is lost.

Rolling Back Transactions

To roll back an entire transaction, or to roll back part of a transaction to a savepoint, use the ROLLBACK statement. For example, either of these statements rolls back the entire current transaction:

```
ROLLBACK WORK;
ROLLBACK;
```

The WORK option of the ROLLBACK statement has no function.

To roll back to a savepoint defined in the current transaction, use the TO option of the ROLLBACK statement. For example, either of these statements rolls back the current transaction to the savepoint named POINT1:

```
SAVEPOINT Point1;
...
ROLLBACK TO SAVEPOINT Point1;
ROLLBACK TO Point1;
```

Defining Transaction Savepoints

To define a savepoint in a transaction, use the SAVEPOINT statement. This statement creates the savepoint named ADD_EMP1 in the current transaction:

```
SAVEPOINT Add_emp1;
```

Creating a savepoint with the same identifier as an earlier savepoint deletes the earlier savepoint. After creating a savepoint, you can roll back to it.

An **active savepoint** is one that was created after the last COMMIT or ROLLBACK statement. The number of active savepoints for each session is unlimited.

Table 1–2 shows a series of SQL statements that illustrates the use of COMMIT, SAVEPOINT, and ROLLBACK statements within a transaction.

Table 1–2 Use of COMMIT, SAVEPOINT, and ROLLBACK

SQL Statement	Results
SAVEPOINT a;	First savepoint of this transaction
DELETE...;	First DML statement of this transaction
SAVEPOINT b;	Second savepoint of this transaction
INSERT INTO...;	Second DML statement of this transaction
SAVEPOINT c;	Third savepoint of this transaction
UPDATE...;	Third DML statement of this transaction
ROLLBACK TO c;	UPDATE statement is rolled back, savepoint C remains defined.
ROLLBACK TO b;	INSERT statement is rolled back, savepoint C is lost, savepoint B remains defined.
ROLLBACK TO c;	ORA-01086
INSERT INTO...;	New DML statement in this transaction
COMMIT;	Commits all actions performed by the first DML statement (the DELETE statement) and the last DML statement (the second INSERT statement). All other statements (the second and the third statements) of the transaction were rolled back before the COMMIT. The savepoint A is no longer active.

Ensuring Repeatable Reads with Read-Only Transactions

By default, Oracle Database guarantees statement-level read consistency, but not transaction-level read consistency. With **statement-level read consistency**, queries in a statement produce consistent data for the duration of the statement, not reflecting changes by other statements. With **transaction-level read consistency** (**repeatable reads**), queries in the transaction produce consistent data for the duration of the transaction, not reflecting changes by other transactions.

To ensure transaction-level read consistency for a transaction that does not include DML statements, specify that the transaction is read-only. The queries in a read-only transaction see only changes committed before the transaction began, so query results are consistent for the duration of the transaction.

A read-only transaction provides transaction-level read consistency without acquiring additional data locks. Therefore, while the read-only transaction is querying data, other transactions can query and update the same data.

A read-only transaction begins with this statement:

```
SET TRANSACTION READ ONLY [ NAME string ];
```

Only DDL statements can precede the `SET TRANSACTION READ ONLY` statement. After the `SET TRANSACTION READ ONLY` statement successfully runs, the transaction can include only `SELECT` (without `FOR UPDATE`), `COMMIT`, `ROLLBACK`, or non-DML statements (such as `SET ROLE`, `ALTER SYSTEM`, and `LOCK TABLE`). A `COMMIT`, `ROLLBACK`, or DDL statement ends the read-only transaction.

> **See Also:** *Oracle Database SQL Language Reference* for more information about the `SET TRANSACTION` statement

Long-running queries sometimes fail because undo information required for consistent read (CR) operations is no longer available. This situation occurs when active transactions overwrite committed undo blocks.

Automatic undo management lets your database administrator (DBA) explicitly control how long the database retains undo information, using the parameter `UNDO_RETENTION`. For example, if `UNDO_RETENTION` is set to 30 minutes, then the database retains all committed undo information for at least 30 minutes, ensuring that all queries running for 30 minutes or less do not encounter the OER error "snapshot too old."

> **See Also:** *Oracle Database Administrator's Guide* for information about long-running queries and resumable space allocation

Using Cursors

PL/SQL implicitly declares a cursor for all SQL data manipulation statements, including queries that return only one row. For queries that return multiple rows, you can explicitly declare a cursor to process the rows individually.

You can think of a **cursor** as a name for a specific private SQL area. A PL/SQL **cursor variable** lets you retrieve multiple rows from a stored subprogram. You can pass cursor variables as parameters in your 3GL application. For more information about cursor variables, see *Oracle Database PL/SQL Language Reference.*

Most Oracle Database users rely on the automatic cursor handling of the database utilities, but programmatic interfaces offer application designers more control over cursors. In application development, a cursor is a named resource available to a program, which can be specifically used for parsing SQL statements embedded within the application.

Topics:

- How Many Cursors Can a Session Have?
- Using a Cursor to Reexecute a Statement
- Scrollable Cursors
- Closing a Cursor
- Canceling a Cursor

How Many Cursors Can a Session Have?

The number of cursors that a session can have open simultaneously is determined by:

- The amount of memory available to the session

- The value of the initialization parameter OPEN_CURSORS, described in *Oracle Database Reference*

Explicitly creating cursors for precompiler programs has advantages in tuning those applications. For example, increasing the number of cursors can reduce the frequency of parsing and improve performance. If you know how many cursors might be required at a given time, you can open that many cursors simultaneously.

Using a Cursor to Reexecute a Statement

After each stage of execution, the cursor retains enough information about the SQL statement to reexecute the statement without starting over, if no other SQL statement was associated with that cursor. The statement can be reexecuted without including the parse stage.

By opening several cursors, the parsed representation of several SQL statements can be saved. Repeated execution of the same SQL statements can thus begin at the describe, define, bind, or run step, saving the repeated cost of opening cursors and parsing.

To understand the performance characteristics of a cursor, the DBA can use the V$SQL dynamic performance view to display the text of the associated query (see *Oracle Database Reference*). Because the results of EXPLAIN PLAN for the original query might differ from the way the database actually processes the query, the DBA can get more precise information by examining these dynamic performance views:

View	Description
V$SQL_PLAN	Execution plan information for each child cursor loaded in the library cache. For more information, see *Oracle Database Reference*.
V$SQL_PLAN_STATISTICS	Execution statistics at the row source level for each child cursor. For more information, see *Oracle Database Reference*.
V$SQL_PLAN_STATISTICS_ALL	Memory usage statistics for row sources that use SQL memory (sort or hash-join). This view concatenates information in V$SQL_PLAN with execution statistics from V$SQL_PLAN_STATISTICS and V$SQL_WORKAREA. For more information, see *Oracle Database Reference*.

Scrollable Cursors

Execution of a cursor puts the results of the query into a set of rows called the result set. A **scrollable cursor** is one whose result rows need not be fetched in forward sequential order. For a scrollable cursor, interfaces exist to fetch previously fetched rows, to fetch the *n*th row in the result set, and to fetch the *n*th row from the current position in the result set.

> **See Also:** *Oracle Call Interface Programmer's Guide* for more information about using scrollable cursors in OCI

Closing a Cursor

Closing a cursor makes its information inaccessible and deallocates the memory that it uses. After a cursor opens, it does not close until the user program terminates its connection to the server.

An OCI program or precompiler application that follows good programming practice explicitly closes open cursors during program execution. If the program terminates without closing an open cursor, then that cursor closes implicitly.

Canceling a Cursor

Canceling a cursor frees resources from the current fetch.The information in the associated private area is lost but the cursor remains open, parsed, and associated with its bind variables.

Note: You cannot cancel cursors using Pro*C/C++ or PL/SQL.

See Also: *Oracle Call Interface Programmer's Guide* for information about canceling a cursor with the `OCIStmtFetch2` statement

Locking Tables Explicitly

Oracle Database has default locking mechanisms that ensure data concurrency, data integrity, and statement-level read consistency. However, you can override these mechanisms by locking tables explicitly. Locking tables explicitly is useful in situations such as these:

- A transaction in your application needs exclusive access to a resource, so that the transaction does not have to wait for other transactions to complete.
- Your application needs transaction-level read consistency (repeatable reads).

  For other ways to ensure transaction-level read consistency, see "Ensuring Repeatable Reads with Read-Only Transactions" on page 1-8) and "Using Serializable Transactions for Concurrency Control" on page 1-24.

To override default locking at the transaction level, use any of these SQL statements:

- `LOCK TABLE` (described in *Oracle Database SQL Language Reference*)
- `SELECT` with the `FOR UPDATE` clause (described in *Oracle Database SQL Language Reference*)
- `SET TRANSACTION` with the `READ ONLY` or `ISOLATION LEVEL SERIALIZABLE` option (described in *Oracle Database SQL Language Reference*)

Locks acquired by these statements are released after the transaction is committed or rolled back.

See Also: *Oracle Database SQL Language Reference* for information about the `ISOLATION_LEVEL` parameter of the `ALTER SESSION` statement

The initialization parameter `DML_LOCKS` (described in *Oracle Database Reference*) determines the maximum number of DML locks. Although its default value is usually enough, you might need to increase it if you use explicit locks.

Caution: If you override the default locking of Oracle Database at any level, ensure that data integrity is guaranteed, data concurrency is acceptable, and deadlocks are either impossible or appropriately handled.

Topics:

- Privileges Required to Acquire Table Locks
- Choosing a Locking Strategy
- Letting Oracle Database Control Table Locking
- Explicitly Acquiring Row Locks
- Examples of Concurrency Under Explicit Locking

Privileges Required to Acquire Table Locks

No special privileges are required to acquire any type of table lock on a table in your own schema. To acquire a table lock on a table in another schema, you must have either the `LOCK ANY TABLE` system privilege or any object privilege (for example, `SELECT` or `UPDATE`) for the table.

Choosing a Locking Strategy

A transaction explicitly acquires the specified table locks when a `LOCK TABLE` statement is executed. A `LOCK TABLE` statement explicitly overrides default locking. When a `LOCK TABLE` statement is issued on a view, the underlying base tables are locked. This statement acquires exclusive table locks for the `employees` and `departments` tables on behalf of the containing transaction:

```
LOCK TABLE employees, departments
   IN EXCLUSIVE MODE NOWAIT;
```

You can specify several tables or views to lock in the same mode; however, only a single lock mode can be specified for each `LOCK TABLE` statement.

> **Note:** When a table is locked, all rows of the table are locked. No other user can modify the table. For information about locking individual rows, see "Explicitly Acquiring Row Locks" on page 1-15.

In the `LOCK TABLE` statement, you can also indicate how long you want to wait for the table lock:

- If you do not want to wait, specify either `NOWAIT` or `WAIT 0`.

  You acquire the table lock only if it is immediately available; otherwise, an error notifies you that the lock is not available now.

- To wait up to *n* seconds to acquire the table lock, specify `WAIT` *n*, where *n* is greater than 0 and less than or equal to 100000.

  If the table lock is still unavailable after *n* seconds, an error notifies you that the lock is not available now.

- To wait indefinitely to acquire the lock, specify neither `NOWAIT` nor `WAIT`.

  The database waits indefinitely until the table is available, locks it, and returns control to you. When the database is running DDL statements concurrently with DML statements, a timeout or deadlock can sometimes result. The database detects such timeouts and deadlocks and returns an error.

Topics:

- When to Lock with ROW SHARE MODE and ROW EXCLUSIVE MODE

- When to Lock with SHARE MODE
- When to Lock with SHARE ROW EXCLUSIVE MODE
- When to Lock with EXCLUSIVE MODE

> **See Also:** *Oracle Database SQL Language Reference* for `LOCK TABLE` statement syntax

When to Lock with ROW SHARE MODE and ROW EXCLUSIVE MODE

`ROW SHARE MODE` and `ROW EXCLUSIVE MODE` table locks offer the highest degree of concurrency. You might use these locks if:

- Your transaction must prevent another transaction from acquiring an intervening share, share row, or exclusive table lock for a table before your transaction can update that table.

  If another transaction acquires an intervening share, share row, or exclusive table lock, no other transactions can update the table until the locking transaction commits or rolls back.
- Your transaction must prevent a table from being altered or dropped before your transaction can modify that table.

When to Lock with SHARE MODE

`SHARE MODE` table locks are rather restrictive data locks. You might use these locks if:

- Your transaction only queries the table, and requires a consistent set of the table data for the duration of the transaction.
- You can hold up other transactions that try to update the locked table, until all transactions that hold `SHARE MODE` locks on the table either commit or roll back.
- Other transactions might acquire concurrent `SHARE MODE` table locks on the same table, also giving them the option of transaction-level read consistency.

> **Caution:** Your transaction might not update the table later in the same transaction. However, if multiple transactions concurrently hold share table locks for the same table, no transaction can update the table (even if row locks are held as the result of a `SELECT FOR UPDATE` statement). Therefore, if concurrent share table locks on the same table are common, updates cannot proceed and deadlocks are common. In this case, use share row exclusive or exclusive table locks instead.

Scenario: Tables `employees` and `budget_tab` require a consistent set of data in a third table, `departments`. For a given department number, you want to update the information in `employees` and `budget_tab`, and ensure that no members are added to the department between these two transactions.

Solution: Lock the `departments` table in `SHARE MODE`, as shown in Example 1–1. Because the `departments` table is rarely updated, locking it probably does not cause many other transactions to wait long.

Example 1–1 LOCK TABLE with SHARE MODE

```
-- Create and populate table:
```

```
DROP TABLE budget_tab;
CREATE TABLE budget_tab (
  sal     NUMBER(8,2),
  deptno  NUMBER(4)
);

INSERT INTO budget_tab (sal, deptno)
  SELECT salary, department_id
  FROM employees;

-- Lock departments and update employees and budget_tab:

LOCK TABLE departments IN SHARE MODE;

UPDATE employees
  SET salary = salary * 1.1
  WHERE department_id IN
    (SELECT department_id FROM departments WHERE location_id = 1700);

UPDATE budget_tab
SET sal = sal * 1.1
WHERE deptno IN
  (SELECT department_id FROM departments WHERE location_id = 1700);

COMMIT;  -- COMMIT releases lock
```

When to Lock with SHARE ROW EXCLUSIVE MODE

You might use a `SHARE ROW EXCLUSIVE MODE` table lock if:

- Your transaction requires both transaction-level read consistency for the specified table and the ability to update the locked table.
- You do not care if other transactions acquire explicit row locks (using `SELECT FOR UPDATE`), which might make `UPDATE` and `INSERT` statements in the locking transaction wait and might cause deadlocks.
- You only want a single transaction to have this action.

When to Lock with EXCLUSIVE MODE

You might use an `EXCLUSIVE MODE` table if:

- Your transaction requires immediate update access to the locked table. When your transaction holds an exclusive table lock, other transactions cannot lock specific rows in the locked table.
- Your transaction also ensures transaction-level read consistency for the locked table until the transaction is committed or rolled back.
- You are not concerned about low levels of data concurrency, making transactions that request exclusive table locks wait in line to update the table sequentially.

Letting Oracle Database Control Table Locking

If you let Oracle Database control table locking, your application needs less programming logic, but also has less control than if you manage the table locks yourself.

Issuing the statement `SET TRANSACTION ISOLATION LEVEL SERIALIZABLE` or `ALTER SESSION ISOLATION LEVEL SERIALIZABLE` preserves ANSI serializability without changing the underlying locking protocol. This technique gives concurrent

access to the table while providing ANSI serializability. Getting table locks greatly reduces concurrency.

See Also:

- *Oracle Database SQL Language Reference* for information about the `SET TRANSACTION` statement
- *Oracle Database SQL Language Reference* for information about the `ALTER SESSION` statements

Change the settings for these parameters only when an instance is shut down. If multiple instances are accessing a single database, then all instances must use the same setting for these parameters.

Explicitly Acquiring Row Locks

You can override default locking with a `SELECT` statement that includes the `FOR UPDATE` clause. This statement acquires exclusive row locks for selected rows (as an `UPDATE` statement does), in anticipation of updating the selected rows in a subsequent statement.

You can use a `SELECT FOR UPDATE` statement to lock a row without actually changing it. For example, several triggers in *Oracle Database PL/SQL Language Reference* show how to implement referential integrity. In the `EMP_DEPT_CHECK` trigger, the row that contains the referenced parent key value is locked to guarantee that it remains for the duration of the transaction; if the parent key is updated or deleted, referential integrity is violated.

`SELECT FOR UPDATE` statements are often used by interactive programs that enable a user to modify fields of one or more specific rows (which might take some time); row locks are acquired so that only a single interactive program user is updating the rows at any given time.

If a `SELECT FOR UPDATE` statement is used when defining a cursor, the rows in the return set are locked when the cursor is opened (before the first fetch) rather than being locked as they are fetched from the cursor. Locks are only released when the transaction that opened the cursor is committed or rolled back, not when the cursor is closed.

Each row in the return set of a `SELECT FOR UPDATE` statement is locked individually; the `SELECT FOR UPDATE` statement waits until the other transaction releases the conflicting row lock. If a `SELECT FOR UPDATE` statement locks many rows in a table, and if the table experiences a lot of update activity, it might be faster to acquire an `EXCLUSIVE` table lock instead.

> **Note:** The return set for a `SELECT FOR UPDATE` might change while the query is running; for example, if columns selected by the query are updated or rows are deleted after the query started. When this happens, `SELECT FOR UPDATE` acquires locks on the rows that did not change, gets a read-consistent snapshot of the table using these locks, and then restarts the query to acquire the remaining locks.
>
> This can cause a deadlock between sessions querying the table concurrently with DML statements when rows are locked in a nonsequential order. To prevent such deadlocks, design your application so that concurrent DML statements on the table do not affect the return set of the query. If this is not feasible, you might want to serialize queries in your application.

By default, the `SELECT FOR UPDATE` statement waits until the requested row lock is acquired. To change this behavior, use the `NOWAIT`, `WAIT`, or `SKIP LOCKED` clause of the `SELECT FOR UPDATE` statement. For information about these clauses, see *Oracle Database SQL Language Reference*.

Examples of Concurrency Under Explicit Locking

Table 1–3 shows how Oracle Database maintains data concurrency, integrity, and consistency when the `LOCK TABLE` statement and the `SELECT` statement with the `FOR UPDATE` clause are used. For brevity, the message text for ORA-00054 ("resource busy and acquire with `NOWAIT` specified") is not included. User-entered text is **bold**.

Table 1–3 Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
`LOCK TABLE hr.departments` **`IN ROW SHARE MODE;`** `Statement processed.`	1	

	2	**`DROP TABLE hr.departments;`**
		
		`DROP TABLE hr.departments`
		`*`
		`ORA-00054`
		
		(Exclusive DDL lock not possible because Transaction 1 has table locked.)

	3	**`LOCK TABLE hr.departments`** **`IN EXCLUSIVE MODE`** **`NOWAIT;`** `ORA-00054`

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
	4	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20** **FOR UPDATE OF location_id;** LOCATION_ID ---------- DALLAS 1 row selected.
UPDATE hr.departments **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** (Waits because Transaction 2 locked same rows.)	5	
	6	**ROLLBACK;** (Releases row locks.)
1 row processed. **ROLLBACK;**	7	
LOCK TABLE hr.departments **IN ROW EXCLUSIVE MODE;** Statement processed.	8	
	9	**LOCK TABLE hr.departments** **IN EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	10	**LOCK TABLE hr.departments** **IN SHARE ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	11	**LOCK TABLE hr.departments** **IN SHARE ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	12	**UPDATE hr.departments** **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** 1 row processed.
	13	**ROLLBACK;**

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
SELECT location_id **FROM hr.departments** **WHERE department_id = 20** **FOR UPDATE OF location_id;** LOCATION_ID ---------- DALLAS 1 row selected.	14	
	15	**UPDATE hr.departments** **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** 1 row processed. (Waits because Transaction 1 locked same rows.)
ROLLBACK;	16	
	17	1 row processed. (Conflicting locks were released.) **ROLLBACK;**
LOCK TABLE hr.departments **IN ROW SHARE MODE** Statement processed.	18	
	19	**LOCK TABLE hr.departments** **IN EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	20	**LOCK TABLE hr.departments** **IN SHARE ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	21	**LOCK TABLE hr.departments** **IN SHARE MODE;** Statement processed.
	22	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20;** LOCATION_ID ---------- DALLAS 1 row selected.

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
	23	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20** **FOR UPDATE OF location_id;** LOCATION_ID ---------- DALLAS 1 row selected.
	24	**UPDATE hr.departments** **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** (Waits because Transaction 1 has conflicting table lock.)
ROLLBACK;	25	
	26	1 row processed. (Conflicting table lock released.) **ROLLBACK;**
LOCK TABLE hr.departments **IN SHARE ROW EXCLUSIVE MODE;** Statement processed.	27	
	28	**LOCK TABLE hr.departments** **IN EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	29	**LOCK TABLE hr.departments** **IN SHARE ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	30	**LOCK TABLE hr.departments** **IN SHARE MODE** **NOWAIT;** ORA-00054
	31	**LOCK TABLE hr.departments** **IN ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	32	**LOCK TABLE hr.departments** **IN SHARE MODE** **NOWAIT;** ORA-00054

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
	33	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20;** LOCATION_ID ---------- DALLAS 1 row selected.
	34	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20** **FOR UPDATE OF location_id;** LOCATION_ID ---------- DALLAS 1 row selected.
	35	**UPDATE hr.departments** **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** (Waits because Transaction 1 has conflicting table lock.)
UPDATE hr.departments **SET location_id = 'NEW YORK'** **WHERE department_id = 20;** (Waits because Transaction 2 locked same rows.)	36	(Deadlock.)
Cancel operation.	37	
ROLLBACK;		
	38	1 row processed.
LOCK TABLE hr.departments **IN EXCLUSIVE MODE;**	39	
	40	**LOCK TABLE hr.departments** **IN EXCLUSIVE MODE;** ORA-00054
	41	**LOCK TABLE hr.departments** **IN ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	42	**LOCK TABLE hr.departments** **IN SHARE MODE;** ORA-00054

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
	43	**LOCK TABLE hr.departments** **IN ROW EXCLUSIVE MODE** **NOWAIT;** ORA-00054
	44	**LOCK TABLE hr.departments** **IN ROW SHARE MODE** **NOWAIT;** ORA-00054
	45	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20;** LOCATION_ID ---------- DALLAS 1 row selected.
	46	**SELECT location_id** **FROM hr.departments** **WHERE department_id = 20** **FOR UPDATE OF location_id;** (Waits because Transaction 1 has conflicting table lock.)
UPDATE hr.departments **SET department_id = 30** **WHERE department_id = 20;** 1 row processed.	47	
COMMIT;	48	
	49	0 rows selected. (Transaction 1 released conflicting lock.)
SET TRANSACTION READ ONLY;	50	
SELECT location_id **FROM hr.departments** **WHERE department_id = 10;** LOCATION_ID ---------- BOSTON	51	
	52	**UPDATE hr.departments** **SET location_id = 'NEW YORK'** **WHERE department_id = 10;** 1 row processed.

Table 1–3 (Cont.) Examples of Concurrency Under Explicit Locking

Transaction 1	Time Point	Transaction 2
SELECT location_id **FROM hr.departments** **WHERE department_id = 10;** LOCATION_ID ---------- BOSTON (Transaction 1 does not see uncommitted data.)	53	
	54	**COMMIT;**
SELECT location_id **FROM hr.departments** **WHERE department_id = 10;** LOCATION_ID ---------- BOSTON (Same result even after Transaction 2 commits.)	55	
COMMIT;	56	
SELECT location_id **FROM hr.departments** **WHERE department_id = 10;** LOCATION_ID ---------- NEW YORK (Sees committed data.)	57	

Using Oracle Lock Management Services (User Locks)

Your applications can use Oracle Lock Management services (user locks) by invoking subprograms the `DBMS_LOCK` package. An application can request a lock of a specific mode, give it a unique name (recognizable in another subprogram in the same or another instance), change the lock mode, and release it. Because a reserved user lock is an Oracle Database lock, it has all the features of a database lock, such as deadlock detection. Ensure that any user locks used in distributed transactions are released upon `COMMIT`, otherwise an undetected deadlock can occur.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for detailed information about the `DBMS_LOCK` package

Topics:

- When to Use User Locks
- Viewing and Monitoring Locks

When to Use User Locks

User locks can help:

- Provide exclusive access to a device, such as a terminal
- Provide application-level enforcement of read locks
- Detect when a lock is released and clean up after the application
- Synchronize applications and enforce sequential processing

Example 1–2 shows how the Pro*COBOL precompiler uses locks to ensure that there are no conflicts when multiple people must access a single device.

Example 1–2 How the Pro*COBOL Precompiler Uses Locks

```
*****************************************************************
* Print Check                                                   *
* Any cashier may issue a refund to a customer returning goods.  *
* Refunds under $50 are given in cash, more than $50 by check.   *
* This code prints the check. One printer is opened by all       *
* the cashiers to avoid the overhead of opening and closing it   *
* for every check, meaning that lines of output from multiple    *
* cashiers can become interleaved if you do not ensure exclusive *
* access to the printer. The DBMS_LOCK package is used to        *
* ensure exclusive access.                                       *
*****************************************************************
CHECK-PRINT
*    Get the lock "handle" for the printer lock.
   MOVE "CHECKPRINT" TO LOCKNAME-ARR.
   MOVE 10 TO LOCKNAME-LEN.
   EXEC SQL EXECUTE
      BEGIN DBMS_LOCK.ALLOCATE_UNIQUE ( :LOCKNAME, :LOCKHANDLE );
      END; END-EXEC.
*   Lock the printer in exclusive mode (default mode).
   EXEC SQL EXECUTE
      BEGIN DBMS_LOCK.REQUEST ( :LOCKHANDLE );
      END; END-EXEC.
*   You now have exclusive use of the printer, print the check.
  ...
*   Unlock the printer so other people can use it
EXEC SQL EXECUTE
      BEGIN DBMS_LOCK.RELEASE ( :LOCKHANDLE );
      END; END-EXEC.
```

Viewing and Monitoring Locks

Table 1–4 describes the Oracle Database facilities that display locking information for ongoing transactions within an instance.

Table 1–4 Ways to Display Locking Information

Tool	Description
Oracle Enterprise Manager Database Control	The Database Locks page shows user locks, all database locks, or locks that are blocking other users or applications. For more information, see *Oracle Database 2 Day + Real Application Clusters Guide*.
Performance Monitoring Data Dictionary Views	See *Oracle Database Administrator's Guide*.

Table 1–4 (Cont.) Ways to Display Locking Information

Tool	Description
UTLLOCKT.SQL	The UTLLOCKT.SQL script displays a simple character lock wait-for graph in tree structured fashion. Using any SQL tool (such as SQL*Plus) to run the script, it prints the sessions in the system that are waiting for locks and the corresponding blocking locks. The location of this script file is operating system dependent. (You must have run the CATBLOCK.SQL script before using UTLLOCKT.SQL.)

Using Serializable Transactions for Concurrency Control

By default, Oracle Database permits concurrently running transactions to modify, add, or delete rows in the same table, and in the same data block. When transaction A changes a table, the changes are invisible to concurrently running transactions until transaction A commits them. If transaction A tries to update or delete a row that transaction B has locked (by issuing a DML or SELECT FOR UPDATE statement), then the DML statement that A issued waits until B either commits or rolls back the transaction. This concurrency model, which provides higher concurrency and thus better performance, is appropriate for most applications.

However, some rare applications require serializable transactions. **Serializable transactions** run concurrently in serialized mode. In **serialized mode**, concurrent transactions can make only the database changes that they could make if they were running serially (that is, one at a time). If a serialized transaction tries to change data that another transaction changed after the serialized transaction began, then error ORA-08177 occurs.

When a serializable transaction fails with ORA-08177, the application can take any of these actions:

- Commit the work executed to that point.
- Run additional, different, statements, perhaps after rolling back to a prior savepoint in the transaction.
- Roll back the transaction and then rerun it.

  The transaction gets a transaction snapshot and the operation is likely to succeed.

  Tip: To minimize the performance overhead of rolling back and re running transactions, put DML statements that might conflict with concurrent transactions near the beginning of the transaction.

  Note: Serializable transactions do not work with deferred segment creation or interval partitioning. Trying to insert data into an empty table with no segment created, or into a partition of an interval partitioned table that does not yet have a segment, causes an error.

Topics:

- Transaction Interaction and Isolation Level
- Setting Isolation Levels
- Serializable Transactions and Referential Integrity
- READ COMMITTED and SERIALIZABLE Isolation Levels

Transaction Interaction and Isolation Level

The ANSI/ISO SQL standard defines three kinds of transaction interaction:

Transaction Interaction	Definition
Dirty read	Transaction A reads uncommitted changes made by transaction B.
Unrepeatable read	Transaction A reads data, transaction B changes the data and commits the changes, and transaction A rereads the data and sees the changes.
Phantom read	Transaction A runs a query, transaction B inserts new rows and commits the change, and transaction A repeats the query and sees the new rows.

The kinds of interactions that a transaction can have is determined by its isolation level. The ANSI/ISO SQL standard defines four transaction isolation levels. Table 1–5 shows what kind of interactions are possible at each isolation level.

Table 1–5 ANSI/ISO SQL Isolation Levels and Possible Transaction Interactions

Isolation Level	Dirty Read	Unrepeatable Read	Phantom Read
`READ UNCOMMITTED`	Possible	Possible	Possible
`READ COMMITTED`	Not possible	Possible	Possible
`REPEATABLE READ`	Not possible	Not possible	Possible
`SERIALIZABLE`	Not possible	Not possible	Not possible

Table 1–6 shows which ANSI/ISO SQL transaction isolation levels Oracle Database provides.

Table 1–6 ANSI/ISO SQL Isolation Levels Provided by Oracle Database

Isolation Level	Provided by Oracle Database
`READ UNCOMMITTED`	No. Oracle Database never permits "dirty reads." Some other database products use this undesirable technique to improve thoughput, but it is not required for high throughput with Oracle Database.
`READ COMMITTED`	Yes, by default. In fact, because an Oracle Database query sees only data that was committed at the beginning of the query (the snapshot time), Oracle Database offers more consistency than the ANSI/ISO SQL standard for `READ COMMITTED` isolation requires.
`REPEATABLE READ`	Yes, if you set the transaction isolation level to `SERIALIZABLE`.
`SERIALIZABLE`	Yes, if you set the transaction isolation level to `SERIALIZABLE`.

Figure 1–1 shows how an arbitrary transaction (that is, one that is either `SERIALIZABLE` or `READ COMMITTED`) interacts with a serializable transaction.

Figure 1–1 Interaction Between Serializable Transaction and Another Transaction

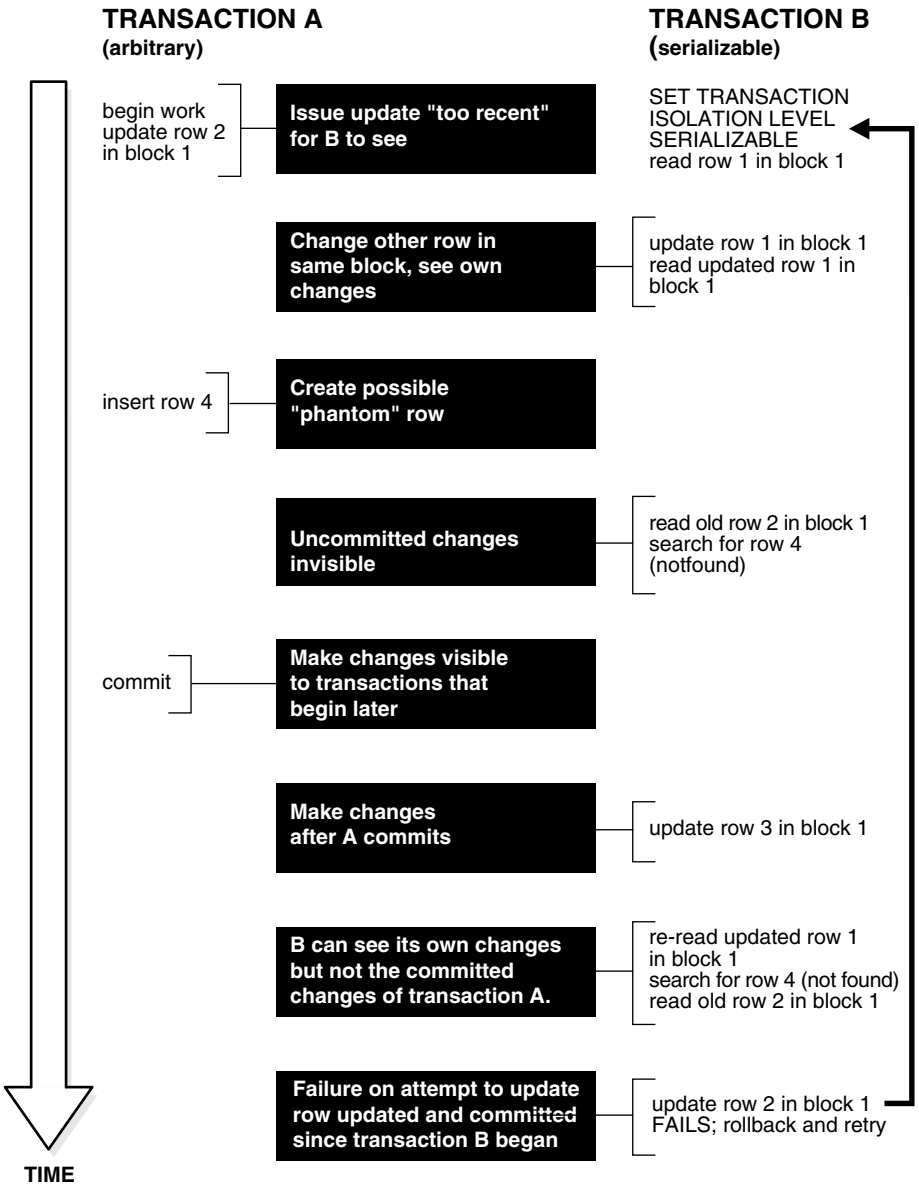

Setting Isolation Levels

To set the transaction isolation level for every transaction in your session, use the `ALTER SESSION` statement, described in *Oracle Database SQL Language Reference*.

To set the transaction isolation level for a specific transaction, use the `ISOLATION LEVEL` clause of the `SET TRANSACTION` statement. The `SET TRANSACTION` statement, described in *Oracle Database SQL Language Reference*, must be the first statement in the transaction.

Note: If you set the transaction isolation level to `SERIALIZABLE`, then you must use the `ALTER TABLE` statement to set the `INITRANS` parameter to at least 3. Use higher values for tables for which many transactions update the same blocks. For more information about `INITRANS`, see *Oracle Database SQL Language Reference*.

Serializable Transactions and Referential Integrity

Because Oracle Database does not use read locks, even in `SERIALIZABLE` transactions, data read by one transaction can be overwritten by another. Therefore, transactions that perform database consistency checks at the application level must not assume that the data they read does not change during the transaction (even though such changes are invisible to the transaction). Code your application-level consistency checks carefully, even when using `SERIALIZABLE` transactions.

In Figure 1–2, transactions A and B (which are either `READ COMMITTED` or `SERIALIZABLE`) perform application-level checks to maintain the referential integrity of the parent/child relationship between two tables. Transaction A queries the parent table to check that it has a row with a specific primary key value before inserting corresponding child rows into the child table. Transaction B queries the child table to check that no child rows exist for a specific primary key value before deleting the corresponding parent row from the parent table. Both transactions assume (but do not ensure) that the data they read does not change before the transaction completes.

Figure 1–2 Referential Integrity Check

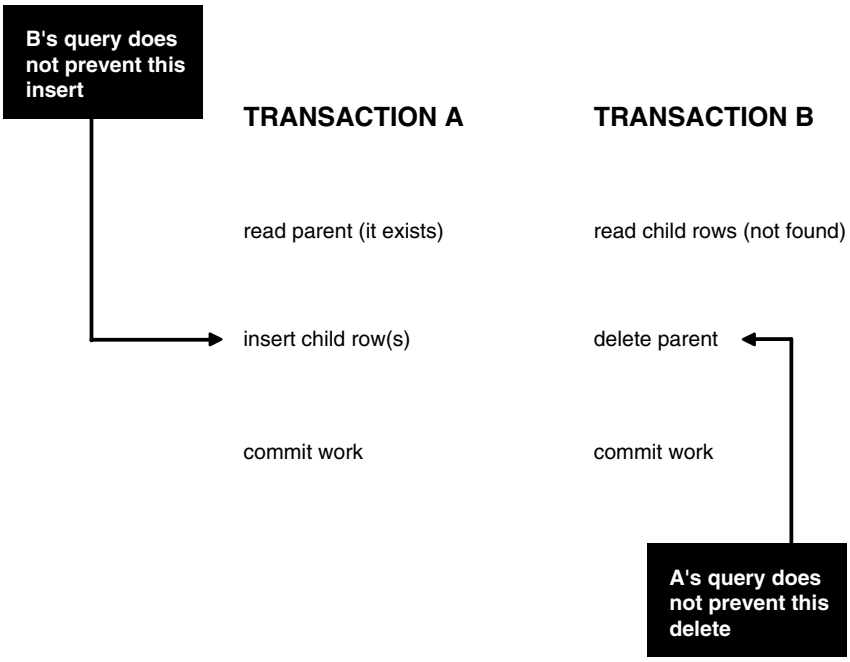

The query by transaction A does not prevent transaction B from deleting the parent row, and the query by transaction B does not prevent transaction A from inserting child rows. Therefore, this can happen:

1. Transaction A queries the parent table and finds the specified parent row.
2. Transaction B queries the child table and finds no child rows for the specified parent row.

3. Having found the specified parent row, transaction A inserts the corresponding child rows into the child table.
4. Having found no child rows for the specified parent row, transaction B deletes the specified parent row from the parent table.

   Now the child rows that transaction A inserted in step 3 have no parent row.

The preceding result can occur even if both A and B are `SERIALIZABLE` transactions, because neither transaction prevents the other from changing the data that it reads to check consistency.

Ensuring that data queried by one transaction is not concurrently changed or deleted by another requires more transaction isolation than the ANSI/ISO SQL standard `SERIALIZABLE` isolation level provides. However, in Oracle Database:

- Transaction A can use a `SELECT FOR UPDATE` statement to query and lock the parent row, thereby preventing transaction B from deleting it.

  For information about the `FOR UPDATE` clause of the `SELECT` statement, see *Oracle Database SQL Language Reference*.
- Transaction B can prevent transaction A from finding the parent row (thereby preventing A from inserting the child rows) by reversing the order of its processing steps. That is, transaction B can:

  1. Delete the parent row.
  2. Query the child table.
  3. If the deleted parent row has child rows in the child table, then roll back the deletion of the parent row.

Alternatively, you can enforce referential integrity with a trigger. Instead of having transaction A query the parent table, define on the child table a row-level `BEFORE INSERT` trigger that does this:

- Queries the parent table with a `SELECT FOR UPDATE` statement, thereby ensuring that if the parent row exists, then it remains in the database for the duration of the transaction that inserts the child rows.
- Rejects the insertion of the child rows if the parent row does not exist.

> **See Also:** *Oracle Database PL/SQL Language Reference* for more information about using triggers to maintain referential integrity between parent and child tables

A trigger runs SQL statements in the context of the triggering statement (that is, the triggering and triggered statements see the database in the same state). Therefore, if a `READ COMMITTED` transaction runs the triggering statement, then the triggered statements see the database as it was when the triggering statement began to execute. If a `SERIALIZABLE` transaction runs the triggering statement, then the triggered statements see the database as it was at the beginning of the transaction. In either case, using `SELECT FOR UPDATE` in the trigger correctly enforces referential integrity.

READ COMMITTED and SERIALIZABLE Isolation Levels

Oracle Database provides two transaction isolation levels, `READ COMMITTED` and `SERIALIZABLE`. Both levels provide a high degree of consistency and concurrency, reduce contention, and are designed for real-world applications. This topic compares them and explains how to choose between them.

Topics:

- Transaction Set Consistency Differences
- Choosing Transaction Isolation Levels

Transaction Set Consistency Differences

An operation (query or transaction) is **transaction set consistent** if all of its read operations return data written by the same set of committed transactions. When an operation is not transaction set consistent, some of its read operations reflect the changes of one set of transactions and others reflect the changes of other sets of transactions; that is, the operation sees the database in a state that reflects no single set of committed transactions.

Topics:

- Oracle Database
- Other Database Systems

Oracle Database Oracle Database transactions with `READ COMMITTED` isolation level are transaction set consistent on an individual-statement basis, because all rows that a query reads must be committed before the query begins.

Oracle Database transactions with `SERIALIZABLE` isolation level are transaction set consistent on an individual-transaction basis, because all statements in a `SERIALIZABLE` transaction run on an image of the database as it was at the beginning of the transaction.

Other Database Systems In other database systems, a single query with `READ UNCOMMITTED` isolation level is not transaction set consistent, because it might see only a subset of the changes made by another transaction. For example, a join of a master table with a detail table can see a master record inserted by another transaction, but not the corresponding details inserted by that transaction (or the reverse). `READ COMMITTED` isolation level avoids this problem, providing more consistency than read-locking systems do.

In read-locking systems, at the cost of preventing concurrent updates, the `REPEATABLE READ` isolation level provides transaction set consistency at the statement level, but not at the transaction level. Due to the absence of phantom read protection, two queries in the same transaction can see data committed by different sets of transactions. In these systems, only the throughput-limiting and deadlock-susceptible `SERIALIZABLE` isolation level provides transaction set consistency at the transaction level.

Choosing Transaction Isolation Levels

The choice of transaction isolation level depends on performance and consistency needs and application coding requirements. There is a trade-off between concurrency (transaction throughput) and consistency. Consider the application and workload when choosing isolation levels for its transactions. Different transactions can have different isolation levels.

For environments with many concurrent users rapidly submitting transactions, consider expected transaction arrival rate, response time demands, and required degree of consistency.

`READ COMMITTED` isolation can provide considerably more concurrency with a somewhat increased risk of inconsistent results (from unrepeatable and phantom reads) for some transactions.

SERIALIZABLE isolation provides somewhat more consistency (by protecting against phantoms and unrepeatable reads), which might be important where a read/write transaction runs a query more than once. However, SERIALIZABLE isolation requires applications to check for the "cannot serialize access" error, and this checking can significantly reduce throughput in an environment with many concurrent transactions accessing the same data for update.

As explained in "Serializable Transactions and Referential Integrity" on page 1-27 reads do not block writes in either READ COMMITTED or SERIALIZABLE transactions.

Table 1–7 summarizes the similarities and differences between READ COMMITTED and SERIALIZABLE transactions.

Table 1–7 Comparison of READ COMMITTED and SERIALIZABLE Transactions

Operation	READ COMMITTED	SERIALIZABLE
Dirty write	Not Possible	Not Possible
Dirty read	Not Possible	Not Possible
Unrepeatable read	Possible	Not Possible
Phantom read	Possible	Not Possible
Compliant with ANSI/ISO SQL 92	Yes	Yes
Read snapshot time	Statement	Transaction
Transaction set consistency	Statement level	Transaction level
Row-level locking	Yes	Yes
Readers block writers	No	No
Writers block readers	No	No
Different-row writers block writers	No	No
Same-row writers block writers	Yes	Yes
Waits for blocking transaction	Yes	Yes
Subject to "cannot serialize access" error	No	Yes
Error after blocking transaction terminates	No	No
Error after blocking transaction commits	No	Yes

Autonomous Transactions

An **autonomous transaction** (AT) is an independent transaction started by another transaction, the **main transaction** (MT). An autonomous transaction lets you suspend the main transaction, do SQL operations, commit or roll back those operations, and then resume the main transaction.

For example, in a stock purchase transaction, you might want to commit customer information regardless of whether the purchase succeeds. Or, you might want to log error messages to a debug table even if the transaction rolls back. Autonomous transactions enable you to do such tasks.

An autonomous transaction runs within an **autonomous scope**; that is, within the scope of an **autonomous routine**—a routine that you mark with the AUTONOMOUS_TRANSACTION pragma. For the definition of **routine** in this context, see *Oracle Database PL/SQL Language Reference*.

Figure 1–3 shows how control flows from the main transaction (MT) to an autonomous transaction (AT) and back again. As you can see, the autonomous transaction can commit multiple transactions (AT1 and AT2) before control returns to the main transaction.

Figure 1–3 Transaction Control Flow

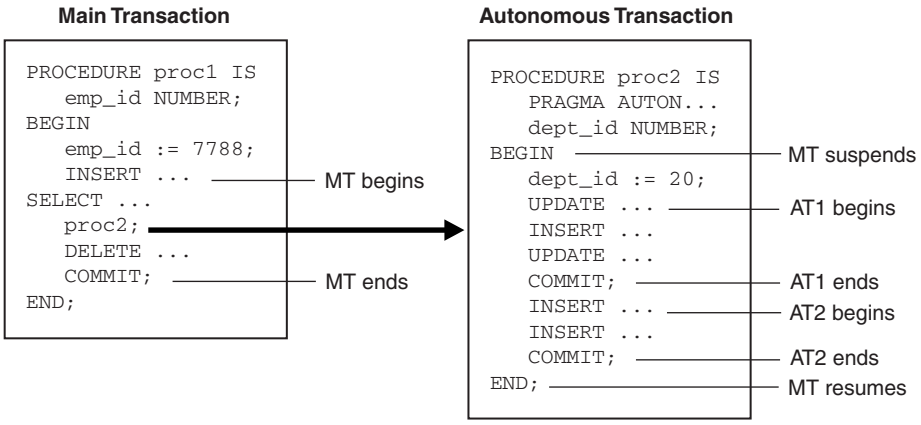

When you enter the executable section of an autonomous transaction, the main transaction suspends. When you exit the transaction, the main transaction resumes. `COMMIT` and `ROLLBACK` end the active autonomous transaction but do not exit the autonomous transaction. As Figure 1–3 shows, when one transaction ends, the next SQL statement begins another transaction.

More characteristics of autonomous transactions:

- The changes an autonomous transaction effects do not depend on the state or the eventual disposition of the main transaction. For example:
  - An autonomous transaction does not see changes made by the main transaction.
  - When an autonomous transaction commits or rolls back, it does not affect the outcome of the main transaction.
- The changes an autonomous transaction effects are visible to other transactions as soon as that autonomous transaction commits. Therefore, users can access the updated information without having to wait for the main transaction to commit.
- Autonomous transactions can start other autonomous transactions.

Figure 1–4 shows some possible sequences that autonomous transactions can follow.

Figure 1–4 Possible Sequences of Autonomous Transactions

A main transaction scope (MT Scope) begins the main transaction, MTx. MTx invokes the first autonomous transaction scope (AT Scope1). MTx suspends. AT Scope 1 begins the transaction Tx1.1.

At Scope 1 commits or rolls back Tx1.1, than ends. MTx resumes.

MTx invokes AT Scope 2. MT suspends, passing control to AT Scope 2 which, initially, is performing queries.

AT Scope 2 then begins Tx2.1 by, say, doing an update. AT Scope 2 commits or rolls back Tx2.1.

Later, AT Scope 2 begins a second transaction, Tx2.2, then commits or rolls it back.

AT Scope 2 performs a few queries, then ends, passing control back to MTx.

MTx invokes AT Scope 3. MTx suspends, AT Scope 3 begins.

AT Scope 3 begins Tx3.1 which, in turn, invokes AT Scope 4. Tx3.1 suspends, AT Scope 4 begins.

AT Scope 4 begins Tx4.1, commits or rolls it back, then ends. AT Scope 3 resumes.

AT Scope 3 commits or rolls back Tx3.1, then ends. MTx resumes.

Finally, MT Scope commits or rolls back MTx, then ends.

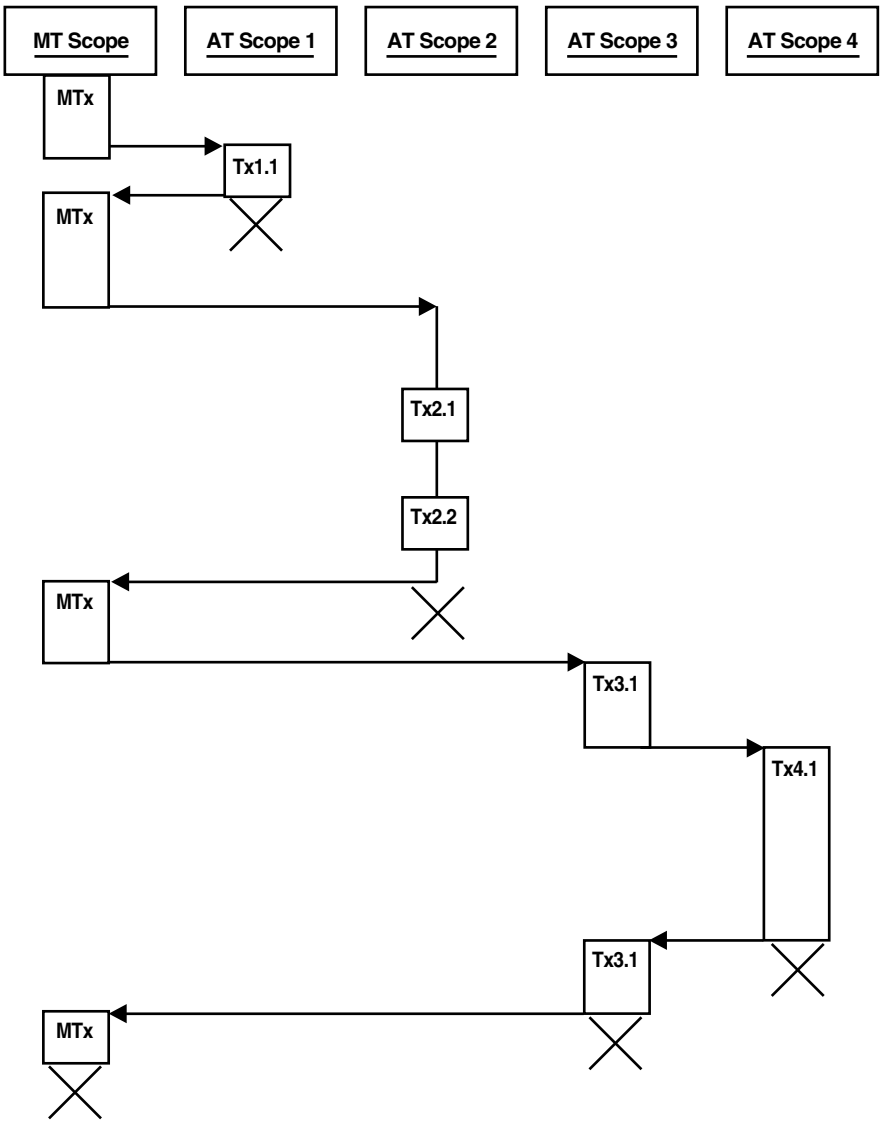

Topics:

- Examples of Autonomous Transactions
- Defining Autonomous Transactions

See Also: *Oracle Database PL/SQL Language Reference* for detailed information about autonomous transactions

Examples of Autonomous Transactions

- Ordering a Product
- Withdrawing Money from a Bank Account

As these examples show, there are four possible outcomes when you use autonomous and main transactions (see Table 1–8). There is no dependency between the outcome of an autonomous transaction and that of a main transaction.

Table 1–8 Possible Transaction Outcomes

Autonomous Transaction	Main Transaction
Commits	Commits
Commits	Rolls back
Rolls back	Commits
Rolls back	Rolls back

Ordering a Product

Figure 1–5 shows an example of a customer ordering a product. The customer information (such as name, address, phone) is committed to a customer information table—even though the sale does not go through.

Figure 1–5 Example: A Buy Order

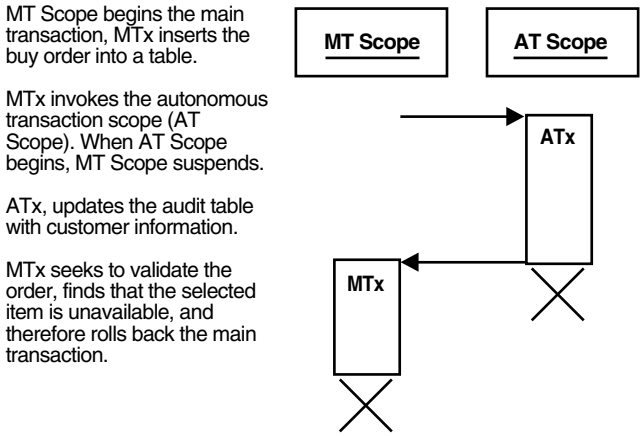

Withdrawing Money from a Bank Account

In this example, a customer tries to withdraw money from a bank account. In the process, a main transaction invokes one of two autonomous transaction scopes (AT Scope 1 or AT Scope 2).

The possible scenarios for this transaction are:

- Scenario 1: Sufficient Funds
- Scenario 2: Insufficient Funds with Overdraft Protection
- Scenario 3: Insufficient Funds Without Overdraft Protection

Scenario 1: Sufficient Funds There are sufficient funds to cover the withdrawal, so the bank releases the funds (see Figure 1–6).

Figure 1–6 Bank Withdrawal—Sufficient Funds

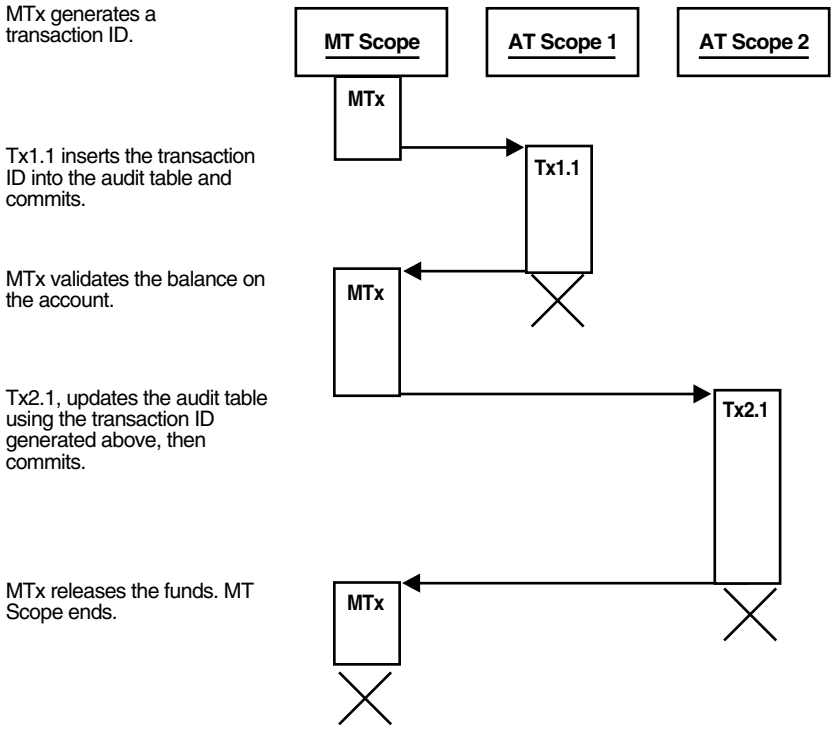

Scenario 2: Insufficient Funds with Overdraft Protection There are insufficient funds to cover the withdrawal, but the customer has overdraft protection, so the bank releases the funds (see Figure 1–7).

Figure 1–7 Bank Withdrawal—Insufficient Funds with Overdraft Protection

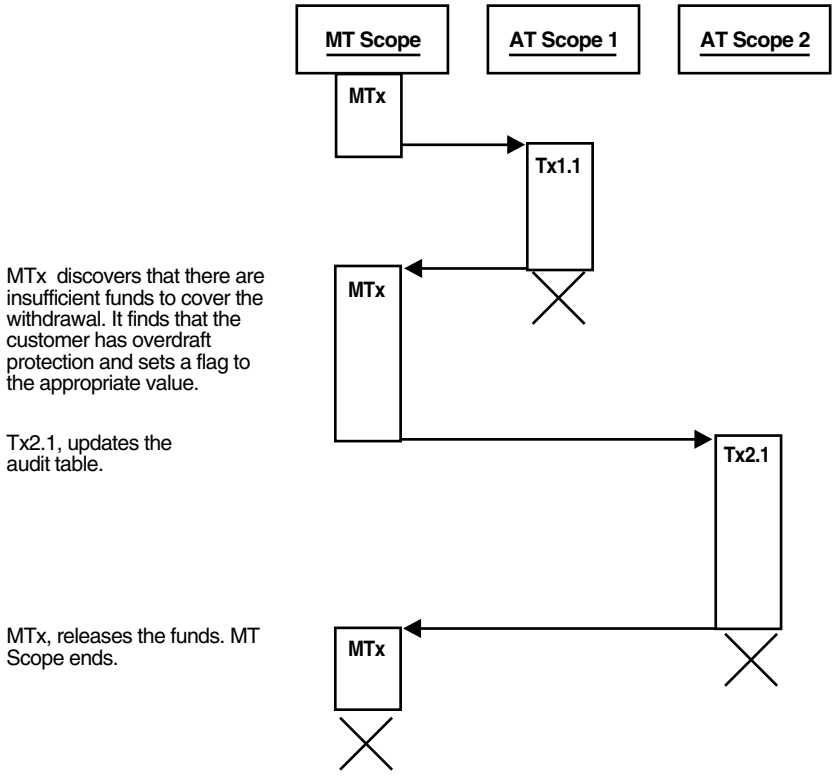

Scenario 3: Insufficient Funds Without Overdraft Protection There are insufficient funds to cover the withdrawal and the customer does not have overdraft protection, so the bank withholds the requested funds (see Figure 1–8).

Figure 1–8 Bank Withdrawal—Insufficient Funds Without Overdraft Protection

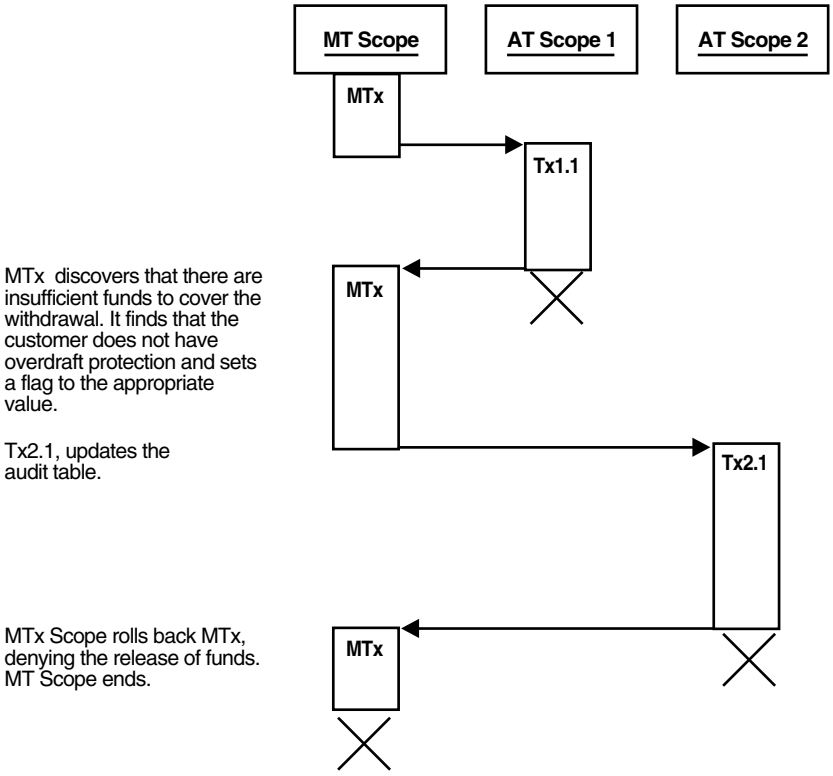

Defining Autonomous Transactions

To define an autonomous transaction, use `PRAGMA AUTONOMOUS_TRANSACTION`, which instructs the PL/SQL compiler to mark the subprogram as autonomous.

In Example 1–3, the function `balance` is autonomous.

Example 1–3 Marking a Package Subprogram as Autonomous

```
-- Create table for package to use:

DROP TABLE accounts;
CREATE TABLE accounts (account INTEGER, balance REAL);

-- Create package:

CREATE OR REPLACE PACKAGE banking AS
  FUNCTION balance (acct_id INTEGER) RETURN REAL;
  -- Additional functions and packages
END banking;
/
CREATE OR REPLACE PACKAGE BODY banking AS
  FUNCTION balance (acct_id INTEGER) RETURN REAL IS
    PRAGMA AUTONOMOUS_TRANSACTION;
    my_bal  REAL;
  BEGIN
    SELECT balance INTO my_bal FROM accounts WHERE account=acct_id;
    RETURN my_bal;
  END;
  -- Additional functions and packages
END banking;
```

/

> **See Also:** *Oracle Database PL/SQL Language Reference* for more information about `PRAGMA AUTONOMOUS_TRANSACTION`

Resuming Execution After Storage Allocation Errors

When a long-running transaction is interrupted by a storage allocation error, the application can suspend the statement that encountered the problem, correct the problem, and then resume executing the statement. This capability, called **resumable storage allocation**, avoids time-consuming rollbacks. It also makes it unnecessary to split the operation into smaller pieces and write code to track its progress.

> **See Also:** *Oracle Database Administrator's Guide* for more information about resumable storage allocation

Topics:

- What Operations Have Resumable Storage Allocation?
- Handling Suspended Storage Allocation

What Operations Have Resumable Storage Allocation?

Queries, DML statements, and some DDL statements have resumable storage allocation after these kinds of errors:

- Out-of-space errors, such as ORA-01653.
- Space-limit errors, such as ORA-01628.
- Space-quota errors, such as ORA-01536.

Resumable storage allocation is possible whether the operation is performed directly by a SQL statement or within SQL*Loader, a stored subprogram, an anonymous PL/SQL block, or an OCI call such as `OCIStmtExecute`.

In dictionary-managed tablespaces, you cannot resume an index- or table-creating operation that encounters the limit for rollback segments or the maximum number of extents. You must use locally managed tablespaces and automatic undo management in combination with resumable storage allocation.

Handling Suspended Storage Allocation

When a statement in an application is suspended because of a storage allocation error, the application does not receive an error code. Therefore, either the application must use an `AFTER SUSPEND` trigger or the DBA must periodically check for suspended statements.

After the problem is corrected (usually by the DBA), the suspended statement automatically resumes execution. If the timeout period expires before the problem is corrected, then the statement raises a `SERVERERROR` exception.

Topics:

- Using an AFTER SUSPEND Trigger in the Application
- Checking for Suspended Statements

Using an AFTER SUSPEND Trigger in the Application

In the application, an AFTER SUSPEND trigger can get information about the problem by invoking subprograms in the DBMS_RESUMABLE package (described in *Oracle Database PL/SQL Packages and Types Reference*). Then the trigger can send the information to an operator, using e-mail (for example).

To reduce the chance of out-of-space errors within the trigger itself, declare the trigger as an autonomous transaction. As an autonomous transaction, the trigger uses a rollback segment in the SYSTEM tablespace. If the trigger encounters a deadlock condition because of locks held by the suspended statement, then the trigger terminates and the application receives the original error code, as if the statement were never suspended. If the trigger encounters an out-of-space condition, then both the trigger and the suspended statement are rolled back. To prevent rollback, use an exception handler in the trigger to wait for the statement to resume.

For general information about triggers, see *Oracle Database PL/SQL Language Reference*.

The trigger in Example 1–4 handles storage errors within the database. For some kinds of errors, the trigger terminates the statement and alerts the DBA, using e-mail. For other errors, which might be temporary, the trigger specifies that the statement waits for eight hours before resuming, expecting the storage problem to be fixed by then. To run this example, you must connect to the database as SYSDBA.

Example 1–4 AFTER SUSPEND Trigger Handles Suspended Storage Allocation

```
-- Create table used by trigger body

DROP TABLE rbs_error;
CREATE TABLE rbs_error (
  SQL_TEXT VARCHAR2(64),
  ERROR_MSG VARCHAR2(64),
  SUSPEND_TIME VARCHAR2(64)
);

-- Resumable Storage Allocation

CREATE OR REPLACE TRIGGER suspend_example
  AFTER SUSPEND
  ON DATABASE
DECLARE
  cur_sid           NUMBER;
  cur_inst          NUMBER;
  err_type          VARCHAR2(64);
  object_owner      VARCHAR2(64);
  object_type       VARCHAR2(64);
  table_space_name  VARCHAR2(64);
  object_name       VARCHAR2(64);
  sub_object_name   VARCHAR2(64);
  msg_body          VARCHAR2(64);
  ret_value         BOOLEAN;
  error_txt         VARCHAR2(64);
  mail_conn         UTL_SMTP.CONNECTION;
BEGIN
 SELECT DISTINCT(SID) INTO cur_sid FROM V$MYSTAT;
 cur_inst := USERENV('instance');
 ret_value := DBMS_RESUMABLE.SPACE_ERROR_INFO
              (err_type,
              object_owner,
              object_type,
              table_space_name,
```

```
              object_name,
              sub_object_name);
 IF object_type = 'ROLLBACK SEGMENT' THEN
   INSERT INTO rbs_error
     (SELECT SQL_TEXT, ERROR_MSG, SUSPEND_TIME
      FROM DBA_RESUMABLE
      WHERE SESSION_ID = cur_sid
      AND INSTANCE_ID = cur_inst);

    SELECT ERROR_MSG INTO error_txt
    FROM DBA_RESUMABLE
    WHERE SESSION_ID = cur_sid
    AND INSTANCE_ID = cur_inst;

    msg_body :=
     'Space error occurred: Space limit reached for rollback segment '
     || object_name || ' on ' || to_char(SYSDATE, 'Month dd, YYYY, HH:MIam')
     || '. Error message was: ' || error_txt;

    mail_conn := UTL_SMTP.OPEN_CONNECTION('localhost', 25);
    UTL_SMTP.HELO(mail_conn, 'localhost');
    UTL_SMTP.MAIL(mail_conn, 'sender@localhost');
    UTL_SMTP.RCPT(mail_conn, 'recipient@localhost');
    UTL_SMTP.DATA(mail_conn, msg_body);
    UTL_SMTP.QUIT(mail_conn);
    DBMS_RESUMABLE.ABORT(cur_sid);
  ELSE
    DBMS_RESUMABLE.SET_TIMEOUT(3600*8);
  END IF;
  COMMIT;
END;
/
```

Checking for Suspended Statements

If the application does not use an `AFTER SUSPEND` trigger, then the DBA must periodically check for suspended statements, using the static data dictionary view `DBA_RESUMABLE` (described in *Oracle Database Reference*).

The DBA can get additional information from the dynamic performance view `V$_SESSION_WAIT` (described in *Oracle Database Reference*).

2

Using SQL Data Types in Database Applications

This chapter explains how to use SQL data types in database applications.

Topics:

- Overview of SQL Data Types
- Representing Character Data
- Representing Numeric Data
- Representing Date and Time Data
- Representing Specialized Data
- Representing Conditional Expressions as Data
- Identifying Rows by Address
- How Oracle Database Converts Data Types
- Metadata for SQL Built-In Functions

See Also:

- *Oracle Database PL/SQL Language Reference* for information about PL/SQL data types
- *Oracle Database PL/SQL Language Reference* for introductory information about Abstract Data Types (ADTs)
- *Oracle Database Object-Relational Developer's Guide* for advanced information about ADTs

  An ADT consists of a data structure and subprograms that manipulate the data. In the static data dictionary view `*_OBJECTS`, the `OBJECT_TYPE` of an ADE is `TYPE`. In the static data dictionary view `*_TYPES`, the `TYPECODE` of an ADE is `OBJECT`.
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for information about LOB data types

  Large object (LOB) data types reference large objects that are stored separately from other data items, such as text, graphic images, video clips, and sound waveforms. LOB data types allow efficient, random, piecewise access to this data.

Overview of SQL Data Types

A data type associates fixed properties with the values that can be inserted in table columns or passed as parameters to subprograms. These properties cause Oracle Database to treat values of different data types differently. For example, Oracle Database can add values of NUMBER data type, but cannot add values of RAW data type.

Oracle Database provides many built-in data types and several categories for user-defined types that can be used as data types.

The Oracle precompilers recognize other data types in embedded SQL programs. These data types are called **external data types** and are associated with host variables. Do not confuse Oracle Database built-in data types and user-defined types with external data types.

See Also:

- *Oracle Database SQL Language Reference* for complete reference information about the SQL data types
- *Oracle Database Concepts* to learn about Oracle Database built-in data types

Representing Character Data

Table 2–1 summarizes the SQL data types that store alphanumeric data.

Table 2–1 SQL Character Data Types

Data Types	Values Stored
CHAR	Fixed-length character literals
NCHAR	Fixed-length Unicode character literals
VARCHAR2	Variable-length character literals
NVARCHAR2	Variable-length Unicode character literals
CLOB, NCLOB	Single-byte and multibyte character strings of up to (4 gigabytes - 1) * (the value obtained from DBMS_LOB.GETCHUNKSIZE)
LONG	Variable-length character strings of up to 2 gigabytes, but with many restrictions This data type is provided only for backward compatibility with existing applications. In general in new applications, use CLOB and NCLOB data types to store large amounts of character data, and BLOB and BFILE to store large amounts of binary data.
LONG RAW	Raw data with up to 2 gigabytes This data type (which is similar to the RAW data type) is provided only for backward compatibility with existing applications.

For a client-server application, if the character set on the client side differs from the character set on the server side, then Oracle Database automatically converts CHAR, VARCHAR2, and LONG data from the database character set (determined by the NLS_LANGUAGE parameter) to the character set defined for the user session.

Topics:

- Specifying Column Lengths as Bytes or Characters
- Choosing Between CHAR and VARCHAR2 Data Types

See Also:

- *Oracle Database SecureFiles and Large Objects Developer's Guide* for information about LOB data types and migration from LONG to LOB data types
- *Oracle Database SQL Language Reference* for restrictions on LONG data types
- *Oracle Database SQL Language Reference* for information about using character literals in SQL statements

Specifying Column Lengths as Bytes or Characters

You can specify the lengths of CHAR and VARCHAR2 columns as either bytes or characters. The lengths of NCHAR and NVARCHAR2 columns are always specified in characters, making them ideal for storing Unicode character literals, where a character might consist of multiple bytes. This table shows some column length specifications and their meanings:

Column Length Specification	Meaning
id VARCHAR2(32 BYTE)	The id column contains up to 32 single-byte characters.
name VARCHAR2(32 CHAR)	The name column contains up to 32 characters of the database character set. If the database character set includes multibyte characters, then the 32 characters can occupy more than 32 bytes.
biography NVARCHAR2(2000)	The biography column contains up to 2000 characters of any Unicode-representable language. The encoding depends on the national character set. The column can contain multibyte values even if the database character set is single-byte.
comment VARCHAR2(2000)	The comment column contains up to 2000 bytes or characters, depending on the value of the initialization parameter NLS_LENGTH_SEMANTICS.

When using a multibyte database character encoding scheme, consider carefully the space required for tables with character columns. If the database character encoding scheme is single-byte, then the number of bytes and the number of characters in a column is the same. If it is multibyte, however, then generally there is no such correspondence. A character might consist of one or more bytes, depending on the specific multibyte encoding scheme and whether shift-in/shift-out control codes are present. To avoid overflowing buffers, specify data as NCHAR or NVARCHAR2 if it might use a Unicode encoding that is different from the database character set.

See Also:

- *Oracle Database Globalization Support Guide* for more information about SQL data types NCHAR and NVARCHAR2
- *Oracle Database SQL Language Reference* for more information about SQL data types NCHAR and NVARCHAR2

Choosing Between CHAR and VARCHAR2 Data Types

When choosing a data type for a column that stores alphanumeric data in a table, consider:

- Space usage

  Oracle Database blank-pads values stored in CHAR columns but not values stored in VARCHAR2 columns. Therefore, VARCHAR2 columns use space more efficiently than CHAR columns.

- Performance

  Because of the blank-padding difference, a full table scan on a large table containing VARCHAR2 columns might read fewer data blocks than a full table scan on a table containing the same data stored in CHAR columns. If your application often performs full table scans on large tables containing character data, then you might be able to improve performance by storing data in VARCHAR2 columns rather than in CHAR columns.

- Comparison semantics

  When you need ANSI compatibility in comparison semantics, use the CHAR data type. When trailing blanks are important in string comparisons, use the VARCHAR2 data type.

  See Also: *Oracle Database SQL Language Reference* for more information about comparison semantics for these data types

- Future compatibility

  The CHAR and VARCHAR2 data types are fully supported. Today, the VARCHAR data type automatically corresponds to the VARCHAR2 data type and is reserved for future use.

Representing Numeric Data

The SQL data types that store numeric data are NUMBER, BINARY_FLOAT, and BINARY_DOUBLE.

The NUMBER data type stores real numbers in either a fixed-point or floating-point format. NUMBER offers up to 38 decimal digits of precision. In a NUMBER column, you can store positive and negative numbers of magnitude 1×10^{-130} through 9.99×10^{125}, and 0. All Oracle Database platforms support NUMBER values.

The BINARY_FLOAT and BINARY_DOUBLE data types store floating-point numbers in the single-precision (32-bit) IEEE 754 format and the double-precision (64-bit) IEEE 754 format, respectively. High-precision values use less space when stored as BINARY_FLOAT and BINARY_DOUBLE than when stored as NUMBER. Arithmetic operations on floating-point numbers are usually faster for BINARY_FLOAT and BINARY_DOUBLE values than for NUMBER values.

In client interfaces that Oracle Database supports, arithmetic operations on BINARY_FLOAT and BINARY_DOUBLE values are performed by the native instruction set that the hardware vendor supplies. The term **native floating-point data type** includes BINARY_FLOAT and BINARY_DOUBLE data types and all implementations of these types in supported client interfaces.

Topics:

- Floating-Point Number Components

- Floating-Point Number Formats
- Comparison Operators for Native Floating-Point Data Types
- Arithmetic Operations with Native Floating-Point Data Types
- Conversion Functions for Floating-Point Data Types
- Client Interfaces for Native Floating-Point Data Types

See Also:

- *Oracle Database SQL Language Reference* for more information about the `NUMBER` data type
- *Oracle Database SQL Language Reference* for more information about the `BINARY_FLOAT` and `BINARY_DOUBLE` data types

Floating-Point Number Components

A floating-point number has these components:

- Binary-valued sign
- Signed exponent
- Significand
- Base

The formula for a floating-point value is:

$(-1)^{sign}.significand.base^{exponent}$

For example, the floating-point value 4.31 is represented:

$(-1)^{0}.431.10^{-2}$

The components of the preceding representation are:

Component Name	Component Value
Sign	0
Significand	431
Base	10
Exponent	-2

Floating-Point Number Formats

A floating-point number format specifies how the components of a floating-point number are represented, thereby determining the range and precision of the values that the format can represent. The **range** is the interval bounded by the smallest and largest values and the **precision** is the number of significant digits. Both range and precision are finite. If a floating-point number is too precise for a given format, then the number is rounded.

How the number is rounded depends on the base of the format, which can be either decimal or binary. A number stored in a decimal floating-point number format (such as `NUMBER`) is rounded to the nearest decimal place (for example, 1000, 10, or 0.01). A number stored in an IEEE 754 format (such as native floating-point data type) is rounded to the nearest binary place (for example, 1024, 512, or 1/64). For calculations that need decimal rounding, use the `NUMBER` data type.

Topics:

- Binary Floating-Point Formats
- Special Values for Native Floating-Point Formats

Binary Floating-Point Formats

This formula determines the value of a floating-point number that uses a binary format:

$$(-1)^{sign}\ 2^{E}\ (bit_0\ bit_1\ bit_2\ \ldots\ bit_{p-1})$$

Table 2–2 describes the components of the preceding formula.

Table 2–2 Binary Floating-Point Format Components

Component	Component Value
`sign`	0 or 1
`E` (exponent)	For single-precision (32-bit) data type, an integer from -126 through 127. For double-precision (64-bit) data type, an integer from -1022 through 1023.
bit_i	0 or 1. (The bit sequence represents a number in base 2.)
`p` (precision)	For single-precision data type, 24. For double-precision data type, 53.

The leading bit of the significand, b_0, must be set (1), except for subnormal numbers (explained later). Therefore, the leading bit is not actually stored, and a binary format provides *n* bits of precision while storing only *n*-1 bits. The IEEE 754 standard defines the in-memory formats for single-precision and double-precision data types, which Table 2–3 shows.

Table 2–3 Summary of Binary Format Storage Parameters

Data Type	Sign Bit	Exponent Bits	Significand Bits	Total Bits
Single-precision	1	8	24 (23 stored)	32
Double-precision	1	11	53 (52 stored)	64

Note: Oracle Database does not support the extended single- and double-precision formats that the IEEE 754 standard defines.

A significand whose leading bit is set is called **normalized**. The IEEE 754 standard defines **subnormal numbers** (also called **denormal numbers**) that are too small to represent with normalized significands. If the significand of a subnormal number were normalized, then its exponent would be too large. Subnormal numbers preserve this property: If *x*-*y*==0.0 (using floating-point subtraction), then *x*==*y*. IEEE 754 formats support subnormal values.

Table 2–4 shows the range and precision of the IEEE 754 single- and double-precision formats and Oracle Database `NUMBER`. Range limits are expressed as positive numbers, but they also apply to absolute values of negative numbers. (The notation "*number* e *exponent*" means *number* * $10^{exponent}$.)

Table 2–4 Range and Precision of Floating-Point Data Types

Range and Precision	Single-precision 32-bit[1]	Double-precision 64-bit[1]	Oracle Database NUMBER Data Type
Maximum positive normal number	3.40282347e+38	1.7976931348623157e+308	< 1.0e126
Minimum positive normal number	1.17549435e-38	2.2250738585072014e-308	1.0e-130
Maximum positive subnormal number	1.17549421e-38	2.2250738585072009e-308	not applicable
Mininum positive subnormal number	1.40129846e-45	4.9406564584124654e-324	not applicable
Precision (decimal digits)	6 - 9	15 - 17	38 - 40

[1] These numbers are from the *IEEE Numerical Computation Guide*.

See Also:

- *Oracle Database SQL Language Reference* for information about literal representation of numeric values
- *Oracle Database SQL Language Reference* for more information about floating-point formats

Special Values for Native Floating-Point Formats

The IEEE 754 standard supports the special values shown in Table 2–5.

Table 2–5 Special Values for Native Floating-Point Formats

Value	Meaning
`+INF`	Positive infinity
`-INF`	Negative infinity
`NaN`	Not a number
`+0`	Positive zero
`-0`	Negative zero

Each value in Table 2–5 is represented by a specific bit pattern, except `NaN`. `NaN`, the result of any undefined operation, is represented by many bit patterns. Some of these bits patterns have the sign bit set and some do not, but the sign bit has no meaning.

The IEEE 754 standard distinguishes between quiet `NaN`s (which do not raise additional exceptions as they propagate through most operations) and signaling `NaN`s (which do). The IEEE 754 standard specifies action for when exceptions are enabled and action for when they are disabled.

In Oracle Database, exceptions cannot be enabled. Oracle Database acts as the IEEE 754 standard specifies for when exceptions are disabled. In particular, Oracle Database does not distinguish between quiet and signaling `NaN`s. You can use Oracle Call Interface (OCI) to retrieve `NaN` values from Oracle Database, but whether a retrieved `NaN` value is signaling or quiet depends on the client platform and is beyond the control of Oracle Database.

The IEEE 754 standard defines these classes of special values:

- Zero
- Subnormal
- Normal
- Infinity
- `NaN`

The values in each class in the preceding list are larger than the values in the classes that precede it in the list (ignoring signs), except `NaN`. `NaN` is unordered with other classes of special values and with itself.

In Oracle Database:

- All `NaN`s are quiet.
- Any non-`NaN` value < `NaN`
- Any `NaN` == any other `NaN`
- All `NaN`s are converted to the same bit pattern.
- -0 is converted to +0.
- IEEE 754 exceptions are not raised.

Comparison Operators for Native Floating-Point Data Types

Oracle Database defines these comparison operators for native floating-point data types:

- Equal to
- Not equal to
- Greater than
- Greater than or equal to
- Less than
- Less than or equal to
- Unordered

Comparisons ignore the sign of zero (`-0` equals `+0`).

> **See Also:** "Special Values for Native Floating-Point Formats" on page 2-7 for more information about comparison results, ordering, and other actions of special values

Arithmetic Operations with Native Floating-Point Data Types

Oracle Database defines these arithmetic operators for native floating-point data types:

- Multiplication
- Division
- Addition
- Subtraction
- Remainder

- Square root

You can define the mode used to round the result of the operation. Exceptions can be raised when operations are performed. Exceptions can also be disabled.

Formerly, Java required floating-point arithmetic to be exactly reproducible. IEEE 754 does not have this requirement. Therefore, results of operations (including arithmetic operations) can be delivered to a destination that uses a range greater than the range that the operands of the operation use.

You can compute the result of a double-precision multiplication at an extended double-precision destination, but the result must be rounded as if the destination were single-precision or double-precision. The range of the result (that is, the number of bits used for the exponent) can use the range supported by the wider (extended double-precision) destination; however, this might cause a double-rounding error in which the least significant bit of the result is incorrect.

This situation can occur only for double-precision multiplication and division on hardware that implements the IA-32 and IA-64 instruction set architecture. Therefore, except for this case, arithmetic for these data types is reproducible across platforms. When the result of a computation is `NaN`, all platforms produce a value for which `IS NAN` is true. However, all platforms do not have to use the same bit pattern.

Conversion Functions for Floating-Point Data Types

Oracle Database defines functions that convert between floating-point and other data types, including string formats that use decimal precision (but precision might be lost during the conversion). For example:

- `TO_BINARY_DOUBLE`, described in *Oracle Database SQL Language Reference*
- `TO_BINARY_FLOAT`, described in *Oracle Database SQL Language Reference*
- `TO_CHAR`, described in *Oracle Database SQL Language Reference*
- `TO_NUMBER`, described in *Oracle Database SQL Language Reference*

Oracle Database can raise exceptions during conversion. The IEEE 754 standard defines these exceptions:

- Invalid
- Inexact
- Divide by zero
- Underflow
- Overflow

Oracle Database does not raise these exceptions for native floating-point data types. Generally, operations that raise exceptions produce the values described in Table 2–6.

Table 2–6 Values Resulting from Exceptions

Exception	Value
Underflow	`0`
Overflow	`-INF`, `+INF`
Invalid Operation	`NaN`
Divide by Zero	`-INF`, `+INF`, `NaN`
Inexact	Any value – rounding was performed

Client Interfaces for Native Floating-Point Data Types

Oracle Database supports native floating-point data types in these client interfaces:

- SQL
- PL/SQL
- Oracle Call Interface (OCI)
- Oracle C++ Call Interface (OCCI)
- Pro*C/C++
- JDBC

Topics:

- OCI Native Floating-Point Data Types SQLT_BFLOAT and SQLT_BDOUBLE
- Native Floating-Point Data Types Supported in ADTs
- Pro*C/C++ Support for Native Floating-Point Data Types

OCI Native Floating-Point Data Types SQLT_BFLOAT and SQLT_BDOUBLE

The OCI API implements the IEEE 754 single- and double-precision native floating-point data types with the data types `SQLT_BFLOAT` and `SQLT_BDOUBLE`, respectively. Conversions between these types and the SQL types `BINARY_FLOAT` and `BINARY_DOUBLE` are exact on platforms that implement the IEEE 754 standard for the C data types `FLOAT` and `DOUBLE`.

See Also: *Oracle Call Interface Programmer's Guide*

Native Floating-Point Data Types Supported in ADTs

Oracle Database supports the SQL data types `BINARY_FLOAT` and `BINARY_DOUBLE` as attributes of ADTs.

Pro*C/C++ Support for Native Floating-Point Data Types

Pro*C/C++ supports the native `FLOAT` and `DOUBLE` data types using the column data types `BINARY_FLOAT` and `BINARY_DOUBLE`. You can use these data types in the same way that Oracle Database `NUMBER` data type is used. You can bind `FLOAT` and `DOUBLE` to `BINARY_FLOAT` and `BINARY_DOUBLE`, respectively, by setting the Pro*C/C++ precompiler command line option `NATIVE_TYPES` to `Y` (yes) when you compile your application.

Representing Date and Time Data

Oracle Database stores date and time data in its own internal format, in fixed-length fields of seven bytes each, corresponding to century, year, month, day, hour, minute, and second.

Table 2–7 summarizes the SQL data types that store date and time data. For more information about these data types, see *Oracle Database SQL Language Reference*.

Table 2–7 SQL Date and Time Data Types

Date Type	Usage
`DATE`	Use to store point-in-time (datetime) values in a table—for example, dates of jobs.

Table 2–7 (Cont.) SQL Date and Time Data Types

Date Type	Usage
TIMESTAMP	Use to store datetime value that are precise to fractional seconds—for example, times of events that must be compared to determine the order in which they occurred.
TIMESTAMP WITH TIME ZONE	Use to store datetime values that must be gathered or coordinated across geographic regions.
TIMESTAMP WITH LOCAL TIME ZONE	Use to store datetime values when the time zone is insignificant—for example, in an application that schedules teleconferences, where participants each see the start and end times for their own time zone. Appropriate for two-tier applications in which you want to display dates and times that use the time zone of the client system. Usually inappropriate for three-tier applications, because data displayed in a Web browser is formatted according to the time zone of the Web server, not the time zone of the browser. The Web server is the database client, so its local time is used.
INTERVAL YEAR TO MONTH	Use to store the difference between two datetime values, where only the year and month are significant—for example, to set a reminder for a date 18 months in the future, or check whether 6 months have elapsed since a particular date.
INTERVAL DAY TO SECOND	Use to store the precise difference between two datetime values—for example, to set a reminder for a time 36 hours in the future or to record the time between the start and end of a race. To represent long spans of time with high precision, use a large number of days.

Topics:

- Displaying Current Date and Time
- Changing the Default Date Format
- Changing the Default Time Format
- Arithmetic Operations with Date and Time Data Types
- Converting Between Date and Time Types
- Importing and Exporting Date and Time Types

See Also: *Oracle Call Interface Programmer's Guide* for a complete description of the Oracle Database internal date format

Displaying Current Date and Time

Use the SQL function SYSDATE to return the system date and time. You can use the FIXED_DATE initialization parameter to set SYSDATE to a constant, which can be useful for testing.

By default, SYSDATE is printed without a BC or AD qualifier. You can add BC to the format string to print the date with the appropriate qualifier, as in Example 2–1.

Example 2–1 Displaying Current Date and Time with AD or BC Qualifier

```
SELECT TO_CHAR(SYSDATE, 'DD-MON-YYYY BC') NOW FROM DUAL;
```

Result:

```
NOW
-----------------------
18-MAR-2009 AD

1 row selected.
```

For input and output of dates, the standard Oracle Database default date format is DD-MON-RR. The RR datetime format element enables you store 20th century dates in the 21st century by specifying only the last two digits of the year. For example, the format '13-NOV-54' refers to the year 1954 in a query issued between 1950 and 2049, but to the year 2054 in a query issued between 2050 and 2099.

> **See Also:** *Oracle Database SQL Language Reference* for information about the RR datetime format element.

Changing the Default Date Format

Use these techniques to change the default date format:

- To change on an instance-wide basis, use the NLS_DATE_FORMAT parameter.
- To change during a session, use the ALTER SESSION statement.

To enter dates that are not in the current default date format, use the TO_DATE function with a format mask, as in Example 2–2.

Example 2–2 Changing the Default Date Format

```
SELECT TO_CHAR(TO_DATE('27-OCT-98', 'DD-MON-RR'), 'YYYY') "Year"
FROM DUAL;
```

Result:

```
Year
----
1998

1 row selected.
```

Be careful when using a date format such as DD-MON-YY. The YY indicates the year in the current century. For example, 31-DEC-92 is December 31, 2092, not 1992 as you might expect. To indicate years in any century other than the current one, use a different format mask, such as the default RR.

Changing the Default Time Format

Time is stored in the 24-hour format: HH24:MI:SS

By default, the time in a DATE column is 12:00:00 A.M. (midnight) if no time portion is specified or if the DATE is truncated.

In a time-only entry, the date portion defaults to the first day of the current month. To enter the time portion of a date, use the TO_DATE function with a format mask indicating the time portion, as in Example 2–3.

Example 2–3 Changing the Default Time Format

```
DROP TABLE birthdays;
CREATE TABLE birthdays (name VARCHAR2(20), day DATE);
INSERT INTO birthdays (name, day)
  VALUES ('ANNIE',
```

```
 TO_DATE('13-NOV-92 10:56 A.M.','DD-MON-YY HH:MI A.M.')
);
```

Arithmetic Operations with Date and Time Data Types

Oracle Database provides features to help with date arithmetic, so that you need not perform your own calculations on the number of seconds in a day, the number of days in each month, and so on. Some useful features include:

- `ADD_MONTHS` function, which returns the date plus the specified number of months.
- `SYSDATE` function, which returns the current date and time set for the operating system on which the database resides.
- `SYSTIMESTAMP` function, which returns the system date, including fractional seconds and time zone, of the system on which the database resides.
- `TRUNC` function, which when applied to a `DATE` value, trims off the time portion so that it represents the very beginning of the day (the stroke of midnight). By truncating two `DATE` values and comparing them, you can determine whether they refer to the same day. You can also use `TRUNC` along with a `GROUP BY` clause to produce daily totals.
- Arithmetic operators such as `+` and `-`. For example, `SYSDATE-7` refers to 7 days before the current system date.
- `INTERVAL` data types, which enable you to represent constants when performing date arithmetic rather than performing your own calculations. For example, you can add or subtract `INTERVAL` constants from `DATE` values or subtract two `DATE` values and compare the result to an `INTERVAL`.
- Comparison operators such as `>`, `<`, `=`, and `BETWEEN`.

Converting Between Date and Time Types

Oracle Database provides several useful functions that enable you to convert to and from datetime data types. Some useful functions include:

- `EXTRACT`, which extracts and returns the value of a specified datetime field from a datetime or interval value expression
- `NUMTODSINTERVAL`, which converts a `NUMBER` or expression that can be implicitly converted to a `NUMBER` value to an `INTERVAL DAY TO SECOND` literal
- `NUMTOYMINTERVAL`, which converts a `NUMBER` or expression that can be implicitly converted to a `NUMBER` value to an `INTERVAL YEAR TO MONTH` literal
- `TO_DATE`, which converts character data to a `DATE` data type
- `TO_CHAR`, which converts `DATE` data to character data
- `TO_DSINTERVAL`, which converts a character string to an `INTERVAL DAY TO SECOND` value
- `TO_TIMESTAMP`, which converts character data to a value of `TIMESTAMP` data type
- `TO_TIMESTAMP_TZ`, which converts character data to a value of `TIMESTAMP WITH TIME ZONE` data type
- `TO_YMINTERVAL`, which converts a character string to an `INTERVAL YEAR TO MONTH` type

> **See Also:** *Oracle Database SQL Language Reference* for details about each function

Importing and Exporting Date and Time Types

`TIMESTAMP WITH TIME ZONE` and `TIMESTAMP WITH LOCAL TIME ZONE` values are always stored in normalized format, so that you can export, import, and compare them without worrying about time zone offsets. `DATE` and `TIMESTAMP` values do not store an associated time zone, and you must adjust them to account for any time zone differences between source and target databases.

Representing Specialized Data

Topics:

- Representing Geographic Data
- Representing Multimedia Data
- Representing Large Amounts of Data
- Representing Searchable Text
- Representing XML
- Representing Dynamically Typed Data
- Representing Data with ANSI/ISO, DB2, and SQL/DS Data Types

Representing Geographic Data

To represent Geographic Information System (GIS) or spatial data in the database, you can use Oracle Spatial features, including the type `MDSYS.SDO_GEOMETRY`. You can store the data in the database by using either an object-relational or a relational model. You can use a set of PL/SQL packages to query and manipulate the data.

> **See Also:** *Oracle Spatial Developer's Guide* to learn how to use `MDSYS.SDO_GEOMETRY`

Representing Multimedia Data

Oracle Multimedia enables Oracle Database to store, manage, and retrieve images, audio, video, or other heterogeneous media data in an integrated fashion with other enterprise information. Oracle Multimedia extends Oracle Database reliability, availability, and data management to multimedia content in traditional, Internet, electronic commerce, and media-rich applications.

Whether you store such multimedia data inside the database as `BLOB` or `BFILE` values, or store it externally on a Web server or other kind of server, you can use Oracle Multimedia to access the data using either an object-relational or a relational model, and manipulate and query the data using a set of ADTs.

Oracle Multimedia provides the `ORDAudio`, `ORDDoc`, `ORDImage`, `ORDImageSignature`, `ORDVideo`, and `SI_StillImage` ADTs (including methods) for these purposes:

- Extracting metadata and attributes from multimedia data
- Retrieving and managing multimedia data from Oracle Multimedia, Web servers, file systems, and other servers

- Performing manipulation operations on image data

  See Also: *Oracle Multimedia Reference* for information about Oracle Multimedia types

Representing Large Amounts of Data

Oracle Database provides several data types for representing large amounts of data. These data types are grouped under the general category of Large Objects (LOBs). Table 2–8 describes the different LOBs.

Table 2–8 Large Object Data Types

Data Type	Name	Description
BLOB	Binary large object	Represents large amounts of binary data such as images, video, or other multimedia data.
CLOB	Character large object	Represents large amounts of character data. CLOB types are stored by using the database character set. Oracle Database stores a CLOB up to 4,000 bytes inline as a VARCHAR2. If the CLOB exceeds this length, then Oracle Database moves the CLOB out of line.
NCLOB	National character large objects	Represents large amounts of character data in National Character Set format.
BFILE	External large object	Stores objects in the operating system's file system, outside of the database files or tablespace. The BFILE type is read-only; other LOB types are read/write. BFILE objects are also sometimes referred to as **external LOBs**.

An instance of type BLOB, CLOB, or NCLOB can exist as either a persistent LOB instance or a temporary LOB instance. Persistent and temporary instances differ as follows:

- A **temporary LOB** instance is declared in the scope of your application.
- A **persistent LOB** instance is created and stored in the database.

Except for declaring, freeing, creating, and committing, operations on persistent and temporary LOB instances are performed the same way.

The RAW and LONG RAW data types store data that is not interpreted by Oracle Database, that is, it is not converted when moving data between different systems. These data types are intended for binary data and byte strings. For example, LONG RAW can store graphics, sound, documents, and arrays of binary data; the interpretation is dependent on the use.

Oracle Net and the Export and Import utilities do not perform character conversion when transmitting RAW or LONG RAW data. When Oracle Database automatically converts RAW or LONG RAW data to and from CHAR data, as is the case when entering RAW data as a literal in an INSERT statement, the database represents the data as one hexadecimal character representing the bit pattern for every four bits of RAW data. For example, one byte of RAW data with bits 11001011 is displayed and entered as CB.

You cannot index LONG RAW data, but you can index RAW data. In earlier releases, the LONG and LONG RAW data types were typically used to store large amounts of data. Use of these types is no longer recommended for development. If your existing application still uses these types, migrate your application to use LOB types. Oracle recommends that you convert LONG RAW columns to binary LOB (BLOB) columns and convert LONG columns to character LOB (CLOB or NCLOB) columns. LOB columns are subject to far fewer restrictions than LONG and LONG RAW columns.

See Also:

- See *Oracle Database SecureFiles and Large Objects Developer's Guide* for more information about LOBs
- See *Oracle Database SQL Language Reference* for restrictions on `LONG` and `LONG RAW` data types

Representing Searchable Text

Rather than writing low-level code to do full-text searches, you can use Oracle Text. It stores the search data in a special kind of index, and lets you query the data with operators and PL/SQL packages. This technology enables you to create your own search engine using data from tables, files, or URLs, and combine the search logic with relational queries. You can also search XML data this way with the XPath notation.

See Also: *Oracle Text Application Developer's Guide* for more information

Representing XML

If you have information stored as files in XML format, or to take an ADT and store it as XML, then you can use the `XMLType` built-in type.

`XMLType` columns store their data as either `CLOB` or binary XML. The `XMLType` constructor can turn an existing object of any data type into an XML object.

When an XML object is inside the database, you can use queries to traverse it (using the XML XPath notation) and extract all or part of its data.

You can also produce XML output from existing relational data and split XML documents across relational tables and columns. You can use these packages to transfer XML data into and out of relational tables:

- `DBMS_XMLQUERY`, which provides database-to-`XMLType` functionality
- `DBMS_XMLGEN`, which converts the results of a SQL query to a canonical XML format
- `DBMS_XMLSAVE`, which provides XML to database-type functionality

You can use these SQL functions to process XML:

- `EXTRACT`, which applies a `VARCHAR2` XPath string and returns an `XMLType` instance containing an XML fragment
- `SYS_XMLAGG`, which aggregates all of the XML documents or fragments represented by an expression and produces a single XML document
- `SYS_XMLGEN`, which takes an expression that evaluates to a particular row and column of the database, and returns an instance of type `XMLType` containing an XML document
- `UPDATEXML`, which takes as arguments an `XMLType` instance and an XPath-value pair and returns an `XMLType` instance with the updated value
- `XMLAGG`, which takes a collection of XML fragments and returns an aggregated XML document
- `XMLCOLATTVAL`, which creates an XML fragment and then expands the resulting XML so that each XML fragment has the name column with the attribute name
- `XMLCONCAT`, which takes as input a series of `XMLType` instances, concatenates the series of elements for each row, and returns the concatenated series

- XMLELEMENT, which takes an element name for identifier, an optional collection of attributes for the element, and arguments that comprise the content of the element
- XMLFOREST, which converts each of its argument parameters to XML, and then returns an XML fragment that is the concatenation of these converted arguments
- XMLSEQUENCE, which either takes as input an XMLType instance and returns a varray of the top-level nodes in the XMLType, or takes as input a REFCURSOR instance, with an optional instance of the XMLFormat object, and returns as an XMLSequence type an XML document for each row of the cursor

  XMLTRANSFORM, which takes as arguments an XMLType instance and an XSL style sheet, applies the style sheet to the instance, and returns an XMLType

See Also:

- *Oracle XML DB Developer's Guide* for details about the XMLType data type
- *Oracle XML Developer's Kit Programmer's Guide* for information about client-side programming with XML
- *Oracle Database SQL Language Reference* for information about XML functions

Representing Dynamically Typed Data

Some languages allow data types to change at run time or let a program check the type of a variable. For example, C has the union keyword and the void * pointer, and Java has the typeof operator and wrapper types such as Number. In Oracle Database, you can create variables and columns that can hold data of any type and test such data values to determine their underlying representation. For example, you can have a single table column represent a numeric value in one row, a string value in another row, and an object in another row.

You can use the built-in ADT SYS.ANYDATA to represent values of any scalar type or ADT. SYS.ANYDATA has methods that accept scalar values of any type, and turn them back into scalars or objects. Similarly, you can use the built-in ADT SYS.ANYDATASET to represent values of any collection type. To check and manipulate type information, use the DBMS_TYPES package, as in Example 2–4. With OCI, use the OCIType, OCIAnyData, and OCIAnyDataSet interfaces.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for information about the DBMS_TYPES package
- *Oracle Database Object-Relational Developer's Guide* for information about the ANYDATA, ANYDATASET, and ANYTYPE types
- *Oracle Call Interface Programmer's Guide* for information about the OCI interfaces

Example 2–4 Accessing Information in a SYS.ANYDATA Column

```
CREATE OR REPLACE TYPE employee_type AS
  OBJECT (empno NUMBER, ename VARCHAR2(10));
/

DROP TABLE mytab;
CREATE TABLE mytab (id NUMBER, data SYS.ANYDATA);
```

```
INSERT INTO mytab (id, data)
VALUES (1, SYS.ANYDATA.ConvertNumber(5));

INSERT INTO mytab (id, data)
VALUES (2, SYS.ANYDATA.ConvertObject(Employee_type(5555, 'john')));

CREATE OR REPLACE PROCEDURE p IS
  CURSOR cur IS SELECT id, data FROM mytab;
  v_id                        mytab.id%TYPE;
  v_data                      mytab.data%TYPE;
  v_type                      SYS.ANYTYPE;
  v_typecode                  PLS_INTEGER;
  v_typename                  VARCHAR2(60);
  v_dummy                     PLS_INTEGER;
  v_n                         NUMBER;
  v_employee                  employee_type;
  non_null_anytype_for_NUMBER exception;
  unknown_typename            exception;
BEGIN
  OPEN cur;
  LOOP
    FETCH cur INTO v_id, v_data;
    EXIT WHEN cur%NOTFOUND;

    /* typecode signifies type represented by v_data.
       GetType also produces a value of type SYS.ANYTYPE with methods you
       can call to find precision and scale of a number, length of a
       string, and so on. */

       v_typecode := v_data.GetType (v_type /* OUT */);

    /* Compare typecode to DBMS_TYPES constants to determine type of data
       and decide how to display it. */

    CASE v_typecode
      WHEN DBMS_TYPES.TYPECODE_NUMBER THEN
        IF v_type IS NOT NULL THEN  -- This condition should never happen.
          RAISE non_null_anytype_for_NUMBER;
        END IF;

      -- For each type, there is a Get method.
      v_dummy := v_data.GetNUMBER (v_n /* OUT */);
      DBMS_OUTPUT.PUT_LINE
        (TO_CHAR(v_id) || ': NUMBER = ' || TO_CHAR(v_n) );

       WHEN DBMS_TYPES.TYPECODE_OBJECT THEN
         v_typename := v_data.GetTypeName();
         IF v_typename NOT IN ('HR.EMPLOYEE_TYPE') THEN
           RAISE unknown_typename;
         END IF;
         v_dummy := v_data.GetObject (v_employee /* OUT */);
         DBMS_OUTPUT.PUT_LINE
           (TO_CHAR(v_id) || ': user-defined type = ' || v_typename ||
            ' ( ' || v_employee.empno || ', ' || v_employee.ename || ' )' );
    END CASE;
  END LOOP;
  CLOSE cur;
EXCEPTION
  WHEN non_null_anytype_for_NUMBER THEN
    RAISE_Application_Error (-20000,
```

```
      'Paradox: the return AnyType instance FROM GetType ' ||
      'should be NULL for all but user-defined types');
  WHEN unknown_typename THEN
    RAISE_Application_Error( -20000, 'Unknown user-defined type ' ||
      v_typename || ' - program written to handle only HR.EMPLOYEE_TYPE');
END;
/

SELECT t.data.gettypename() AS "Type Name" FROM mytab t;
```

Result:

```
Type Name
--------------------------------------------------------------------------------
SYS.NUMBER
HR.EMPLOYEE_TYPE

2 rows selected.
```

Representing Data with ANSI/ISO, DB2, and SQL/DS Data Types

You can define columns of tables in Oracle Database through ANSI/ISO, DB2, and SQL/DS data types. Oracle Database internally converts such data types to Oracle Database data types.

The ANSI data type conversions are shown in Table 2–9. The ANSI/ISO data types `NUMERIC`, `DECIMAL`, and `DEC` can specify only fixed-point numbers. For these data types, `s` defaults to 0.

Table 2–9 ANSI Data Type Conversions to Oracle Database Data Types

ANSI SQL Data Type	Oracle Database Data Type
CHARACTER (*n*)	CHAR (*n*)
CHAR (*n*)	
NUMERIC (*p*,*s*)	NUMBER (*p*,*s*)
DECIMAL (*p*,*s*)	
DEC (*p*,*s*)	
INTEGER	NUMBER (38)
INT	
SMALLINT	
FLOAT (*p*)	FLOAT (*p*)
REAL	FLOAT (63)
DOUBLE PRECISION	FLOAT (126)
CHARACTER VARYING(*n*)	VARCHAR2 (*n*)
CHAR VARYING(*n*)	
TIMESTAMP	TIMESTAMP
TIMESTAMP WITH TIME ZONE	TIMESTAMP WITH TIME ZONE

Table 2–10 shows the SQL/DS and DB2 conversions.

Table 2–10 SQL/DS, DB2 Data Type Conversions to Oracle Database Data Types

DB2 or SQL/DS Data Type	Oracle Database Data Type
CHARACTER (*n*)	CHAR (*n*)
VARCHAR (*n*)	VARCHAR2 (*n*)
LONG VARCHAR	LONG
DECIMAL (*p*,*s*)	NUMBER (*p*,*s*)
INTEGER	NUMBER (38)
SMALLINT	
FLOAT (*p*)	FLOAT (*p*)
DATE	DATE
TIMESTAMP	TIMESTAMP

The data types TIME, GRAPHIC, VARGRAPHIC, and LONG VARGRAPHIC of IBM products SQL/DS and DB2 have no corresponding Oracle Database data type, and they cannot be used.

Representing Conditional Expressions as Data

The Oracle Expression Filter feature enables you to store conditional expressions as data in the database. The Oracle Expression Filter provides a mechanism that you can use to place a constraint on a VARCHAR2 column to ensure that the values stored are valid SQL WHERE clause expressions. This mechanism also identifies the set of attributes that are legal to reference in the conditional expressions.

Scenario: You created the following table, in which each row holds data for a stock-trading account holder, and you want to define a column that stores information about the stocks in which each trader is interested as a conditional expression.

```
DROP TABLE traders;
CREATE TABLE traders (
  name     VARCHAR2(10),
  email    VARCHAR2(20),
  interest VARCHAR2(30)
);
```

Solution:

1. Create a type with attributes for the trading symbol, limit price, and amount of change in the stock price:

   ```
   CREATE OR REPLACE TYPE ticker AS OBJECT (
     symbol VARCHAR2(20),
     price  NUMBER,
     change NUMBER
   );
   /
   ```

2. Create an attribute set based on the type ticker:

   ```
   BEGIN
     DBMS_EXPFIL.DROP_ATTRIBUTE_SET (attr_set  => 'ticker');
   END;
   /
   BEGIN
   ```

```
  DBMS_EXPFIL.CREATE_ATTRIBUTE_SET
  (attr_set  => 'ticker',
  from_type => 'YES');
END;
/
```

3. Associate the attribute set with the expression set stored in the column `trader.interest`:

```
BEGIN
  DBMS_EXPFIL.ASSIGN_ATTRIBUTE_SET
  (attr_set => 'ticker',
  expr_tab => 'traders',
  expr_col => 'interest');
  END;
/
```

   The preceding code ensures that the `interest` column stores valid conditional expressions.

4. Populate the table with trader names, e-mail addresses, and conditional expressions that represent stocks in which the trader is interested, at particular prices. For example:

```
INSERT INTO traders (name, email, interest)
VALUES ('Vishu', 'vishu@example.com', 'symbol = ''ABC'' AND price > 25');
```

5. Use the `EVALUATE` operator to identify the conditional expressions that evaluate to `TRUE` for a given data item. For example, this query returns traders who are interested in the stock quote `(symbol='ABC', price=31, change=5.2)`:

```
SELECT name, email
FROM traders
WHERE EVALUATE (
  interest,
  'symbol=>''ABC'',
  price=>31,
  change=>5.2'
) = 1;
```

Result:

```
NAME       EMAIL
---------- --------------------
Vishu      vishu@example.com

1 row selected.
```

To speed up this type of query, you can create an Oracle Expression Filter index on the `interest` column.

> **See Also:** *Oracle Database Rules Manager and Expression Filter Developer's Guide* for details on Oracle Expression Filter

Identifying Rows by Address

The fastest way to access a row is by its address, or **rowid**, which uniquely identifies it. Different rows in the same data block can have the same rowid only if they are in different clustered tables. If a row is larger than one data block, then its rowid identifies its initial row piece.

To see rowids, you query the ROWID pseudocolumn, whose value is a string that represents the address of the row. The string has the data type ROWID or UROWID.

> **See Also:** *Oracle Database SQL Language Reference* for more information about the ROWID pseudocolumn

Topics:

- Querying the ROWID Pseudocolumn
- ROWID Data Type
- UROWID Data Type

Querying the ROWID Pseudocolumn

Each table in Oracle Database has a pseudocolumn named ROWID, which can appear in a query in either the SELECT list or the WHERE clause.

Example 2–5 uses the ROWID pseudocolumn in the SELECT list of a query. The rowids show how the rows of the table are stored.

Example 2–5 Querying the ROWID Pseudocolumn

```
DROP TABLE t_tab;  -- in case it exists
CREATE TABLE t_tab (col1 ROWID);

INSERT INTO t_tab (col1)
SELECT ROWID
FROM employees
WHERE employee_id > 199;
```

Query:

```
SELECT employee_id, rowid
FROM employees
WHERE employee_id > 199;
```

ROWID varies, but result is similar to:

```
EMPLOYEE_ID ROWID
----------- ------------------
        200 AAAPeSAAFAAAABTAAC
        201 AAAPeSAAFAAAABTAAD
        202 AAAPeSAAFAAAABTAAE
        203 AAAPeSAAFAAAABTAAF
        204 AAAPeSAAFAAAABTAAG
        205 AAAPeSAAFAAAABTAAH
        206 AAAPeSAAFAAAABTAAI

7 rows selected.
```

Query:

```
SELECT * FROM t_tab;
```

COL1 varies, but result is similar to:

```
COL1
------------------
AAAPeSAAFAAAABTAAC
AAAPeSAAFAAAABTAAD
AAAPeSAAFAAAABTAAE
```

```
AAAPeSAAFAAAABTAAF
AAAPeSAAFAAAABTAAG
AAAPeSAAFAAAABTAAH
AAAPeSAAFAAAABTAAI

7 rows selected.
```

ROWID Data Type

In tables that are not index-organized, and in foreign tables, the values of the ROWID pseudocolumn have the data type ROWID. The format of this data type is either restricted, extended or external binary.

> **Note:** You can create tables and clusters that have columns of the type ROWID, but the values of these columns are not guaranteed to be valid rowids.

Topics:

- Restricted ROWID
- Extended ROWID
- External Binary ROWID

Restricted ROWID

Internally, the ROWID is a structure that holds information that the database server must access a row. The restricted internal ROWID is 6 bytes on most platforms. Each restricted rowid includes these data:

- Data file identifier
- Block identifier
- Row identifier

The restricted ROWID pseudocolumn is returned to client applications in the form of an 18-character string with a hexadecimal encoding of the data block, row, and data file components of the ROWID.

Extended ROWID

The extended ROWID data type includes the data in the restricted rowid plus a data object number. The data object number is an identification number assigned to every database segment. The extended internal ROWID is 10 bytes on most platforms.

Data in an extended ROWID pseudocolumn is returned to the client application in the form of an 18-character string (for example, "AAAA8mAALAAAAQkAAA"), which represents a base 64 encoding of the components of the extended ROWID in a four-piece format, OOOOOOFFFBBBBBBRRR. Extended rowids are not available directly. You can use a supplied package, DBMS_ROWID, to interpret extended rowid contents. The package functions extract and provide information that is available directly from a restricted rowid and information specific to extended rowids.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for information about the DBMS_ROWID package

External Binary ROWID

Some client applications use a binary form of the ROWID. For example, OCI and some precompiler applications can map the ROWID data type to a 3GL structure on bind or define calls. The size of the binary ROWID is the same for extended and restricted ROWIDs. The information for the extended ROWID is included in an unused field of the restricted ROWID structure.

The format of the extended binary ROWID, expressed as a C struct, is as follows:

```
struct riddef {
    ub4    ridobjnum; /* data obj#--this field is
                         unused in restricted ROWIDs */
    ub2    ridfilenum;
    ub1    filler;
    ub4    ridblocknum;
    ub2    ridslotnum;
}
```

UROWID Data Type

The rows of some tables have addresses that are not physical or permanent or were not generated by Oracle Database. For example, the row addresses of index-organized tables are stored in index leaves, which can move. Oracle Database provides these tables with logical row identifiers, called **logical rowids**. Rowids of foreign tables, such as DB2 tables accessed through a gateway, are not standard Oracle Database rowids. Oracle Database provides foreign tables with identifiers called **foreign rowids**.Oracle Database uses universal rowids (urowids) to store the addresses of index-organized and foreign tables. Both types of urowid are stored in the ROWID pseudocolumn, as are the physical rowids of heap-organized tables.Oracle Database creates logical rowids based on the primary key of the table. The logical rowids do not change if the primary key does not change. The ROWID pseudocolumn of an index-organized table has a data type of UROWID. You can access this pseudocolumn as you would access the ROWID pseudocolumn of a heap-organized table (that is, using a SELECT ROWID statement). To store the rowids of an index-organized table, define a column of type UROWID for the table and retrieve the value of the ROWID pseudocolumn into that column.

How Oracle Database Converts Data Types

In some cases, Oracle Database accepts data of one data type where it expects data of a different data type. Generally, an expression cannot contain values with different data types. However, Oracle Database can use various SQL functions to automatically convert data to the expected data type.

> **See Also:** *Oracle Database SQL Language Reference* for details about data type conversion

Topics:

- Data Type Conversion During Assignments
- Data Type Conversion During Expression Evaluation

Data Type Conversion During Assignments

The data type conversion for an assignment succeeds if Oracle Database can convert the data type of the value to be assigned to the data type of the target.

Assume that `test_package`, its public variable `var1`, and `table1_tab` are declared as follows:

```
CREATE OR REPLACE PACKAGE test_package AS
  var1 CHAR(5);
END;
/

DROP TABLE table1_tab;
CREATE TABLE table1_tab (col1 NUMBER);
```

In the assignment

```
variable := expression
```

the data type of *`expression`* must be either the same as, or convertible to, the data type of *`variable`*. For example, for this assignment, Oracle Database automatically converts zero to the data type of `var1`, which is `CHAR(5)`:

```
var1 := 0;
```

In the statement

```
INSERT INTO table1_tab (col1) VALUES (expression)
```

the data type of *`expression`* must be either the same as, or convertible to, the data types of `col1`. For example, for this statement, Oracle Database automatically converts the string '19' to the data type of `col1`, which is `NUMBER`:

```
INSERT INTO table1_tab (col1) VALUES ('19')
```

In the statement

```
UPDATE table1_tab SET column = expression
```

the data type of *`expression`* must be either the same as, or convertible to, the data type of *`column`*. For example, for this statement, Oracle Database automatically converts the string '30' to the data type of `col1`, which is `NUMBER`:

```
UPDATE table1_tab SET col1 = '30';
```

In the statement

```
SELECT column INTO variable FROM table1_tab
```

the data type of *`column`* must be either the same as, or convertible to, the data type of *`variable`*. For example, for this statement, Oracle Database automatically converts the value selected from `col1`, which is 30, to the data type of `var1`, which is `CHAR(5)`:

```
SELECT col1 INTO var1 FROM table1_tab WHERE col1 = 30;
```

Data Type Conversion During Expression Evaluation

For expression evaluation, Oracle Database can automatically perform the same conversions as for assignments. An expression is converted to a type based on its context. For example, operands to arithmetic operators are converted to `NUMBER`, and operands to string functions are converted to `VARCHAR2`.

Oracle Database can automatically convert:

- `VARCHAR2` or `CHAR` to `NUMBER`

- VARCHAR2 or CHAR to DATE

Character to NUMBER conversions succeed only if the character string represents a valid number. Character to DATE conversions succeed only if the character string satisfies the session default format, which is specified by the initialization parameter NLS_DATE_FORMAT.

Some common types of expressions are:

- Simple expressions, such as:

```
commission + '500'
```

- Boolean expressions, such as:

```
bonus > salary / '10'
```

- Subprogram calls, such as:

```
MOD (counter, '2')
```

- WHERE clause conditions, such as:

```
WHERE hiredate = TO_DATE('1997-01-01','yyyy-mm-dd')
```

- WHERE clause conditions, such as:

```
WHERE rowid = 'AAAAaoAATAAAADAAA'
```

In general, Oracle Database uses the rule for expression evaluation when a data type conversion is needed in places not covered by the rule for assignment conversions.

In assignments of the form:

```
variable := expression
```

Oracle Database first evaluates *expression* using the conversion rules for expressions; *expression* can be as simple or complex as desired. If it succeeds, then the evaluation of *expression* results in a single value and data type. Then, Oracle Database tries to assign this value to the target variable using the conversion rules for assignments.

Metadata for SQL Built-In Functions

You can see metadata for SQL built-in functions with the dynamic performance views V$SQLFN_METADATA (which has general metadata) and V$SQLFN_ARG_METADATA (which has metadata about arguments). You can join these views on the column FUNCID. For functions with unlimited arguments, such as LEAST and GREATEST, V$SQLFN_ARG_METADATA has only one row for each repeating argument.

These views enable third-party tools to leverage SQL built-in functions without maintaining their metadata in the application layer.

> **See Also:** *Oracle Database Reference* for detailed information about the dynamic performance views V$SQLFN_METADATA and V$SQLFN_ARG_METADATA

Often, an argument for a SQL built-in function can have any data type in a data type family. Table 2–11 shows which data types belong to which families.

Table 2–11 Data Type Families

Family	Data Types
STRING	CHARACTER
	VARCHAR2
	CLOB
	NCHAR
	NVARCHAR2
	NCLOB
NUMERIC	NUMBER
	BINARY_FLOAT
	BINARY_DOUBLE
DATETYPE	DATE
	TIMESTAMP
	TIMESTAMP WITH TIME ZONE
	TIMESTAMP WITH LOCAL TIME ZONE
	INTERVAL YEAR TO MONTH
	INTERVAL DAY TO SECOND
BINARY	BLOB
	RAW
	LONGRAW

ARG*n* Data Type

In the view V$SQLFN_METADATA, ARG*n* is the data type of a function whose return value has the same data type as its *n*th argument. For example:

- The MAX function returns a value that has the data type of its first argument, so the MAX function has data type ARG1.
- The DECODE function returns a value that has the data type of its third argument, so the DECODE function has data type ARG3.

EXPR Data Type

In the view V$SQLFN_ARG_METADATA, EXPR is the data type of an argument that can be any expression. An expression is either a single value or a combination of values and SQL functions that has a single value.

Table 2–12 Display Types of SQL Built-In Functions

Display Type	Description	Example
NORMAL	FUNC(A,B,...)	LEAST(A,B,C)
ARITHMETIC	A FUNC B)	A+B
PARENTHESIS	FUNC()	SYS_GUID()
RELOP	A FUNC B)	A IN B
CASE_LIKE	CASE statement or DECODE decode	
NOPAREN	FUNC	SYSDATE

3

Using Regular Expressions in Database Applications

This chapter explains how to use regular expressions in database applications.

Topics:

- Overview of Regular Expressions
- Metacharacters in Regular Expressions
- Using Regular Expressions in SQL Statements: Scenarios

See Also:

- *Oracle Database SQL Language Reference* for information about Oracle Database SQL functions for regular expressions
- *Oracle Database Globalization Support Guide* for details on using SQL regular expression functions in a multilingual environment
- *Oracle Regular Expressions Pocket Reference* by Jonathan Gennick, O'Reilly & Associates
- *Mastering Regular Expressions* by Jeffrey E. F. Friedl, O'Reilly & Associates

Overview of Regular Expressions

Topics:

- What Are Regular Expressions?
- How Are Regular Expressions Useful?
- Oracle Database Implementation of Regular Expressions
- Oracle Database Support for the POSIX Regular Expression Standard

What Are Regular Expressions?

Regular expressions enable you to search for patterns in string data by using standardized syntax conventions. You specify a regular expression through these types of characters:

- Metacharacters, which are operators that specify search algorithms
- Literals, which are the characters for which you are searching

A regular expression can specify complex patterns of character sequences. For example, this regular expression searches for the literals `f` or `ht`, the `t` literal, the `p` literal optionally followed by the `s` literal, and finally the colon (`:`) literal:

```
(f|ht)tps?:
```

The parentheses are metacharacters that group a series of pattern elements to a single element; the pipe symbol (`|`) matches an alternative in the group. The question mark (`?`) is a metacharacter indicating that the preceding pattern, in this case the `s` character, is optional. Thus, the preceding regular expression matches the `http:`, `https:`, `ftp:`, and `ftps:` strings.

How Are Regular Expressions Useful?

Regular expressions are a powerful text processing component of programming languages such as PERL and Java. For example, a PERL script can process each HTML file in a directory, read its contents into a scalar variable as a single string, and then use regular expressions to search for URLs in the string. One reason that many developers write in PERL is for its robust pattern matching functionality.

Oracle Database support of regular expressions enables developers to implement complex match logic in the database. This technique is useful for these reasons:

- By centralizing match logic in Oracle Database, you avoid intensive string processing of SQL results sets by middle-tier applications. For example, life science customers often rely on PERL to do pattern analysis on bioinformatics data stored in huge databases of DNA and proteins. Previously, finding a match for a protein sequence such as `[AG].{4}GK[ST]` was handled in the middle tier. The SQL regular expression functions move the processing logic closer to the data, thereby providing a more efficient solution.
- Before Oracle Database 10*g*, developers often coded data validation logic on the client, requiring the same validation logic to be duplicated for multiple clients. Using server-side regular expressions to enforce constraints solves this problem.
- The built-in SQL and PL/SQL regular expression functions and conditions make string manipulations more powerful and less cumbersome than in previous releases of Oracle Database.

Oracle Database Implementation of Regular Expressions

Oracle Database implements regular expression support with a set of Oracle Database SQL functions and conditions that enable you to search and manipulate string data. You can use these functions in any environment that supports Oracle Database SQL. You can use these functions on a text literal, bind variable, or any column that holds character data such as `CHAR`, `NCHAR`, `CLOB`, `NCLOB`, `NVARCHAR2`, and `VARCHAR2` (but not `LONG`).

Table 3–1 describes the regular expression functions and conditions.

Table 3–1 SQL Regular Expression Functions and Conditions

SQL Element	Category	Description
REGEXP_LIKE	Condition	Searches a character column for a pattern. Use this function in the WHERE clause of a query to return rows matching a regular expression. The condition is also valid in a constraint or as a PL/SQL function returning a boolean. This WHERE clause filters employees with a first name of Steven or Stephen: `WHERE REGEXP_LIKE(first_name, '^Ste(v|ph)en$')`
REGEXP_REPLACE	Function	Searches for a pattern in a character column and replaces each occurrence of that pattern with the specified string. These function call puts a space after each character in the country_name column: `REGEXP_REPLACE(country_name, '(.)', '\1 ')`
REGEXP_INSTR	Function	Searches a string or substring for a given occurrence of a regular expression pattern (a substring) and returns an integer indicating the position in the string or substring where the match is found. You specify which occurrence you want to find and the start position. This function call performs a boolean test for a valid e-mail address in the email column: `REGEXP_INSTR(email, '\w+@\w+(\.\w+)+') > 0`
REGEXP_SUBSTR	Function	Searches a string or substring for a given occurrence of a regular expression pattern (a substring) and returns the substring itself. You specify which occurrence you want to find and the start position. This function call uses the x flag to match the first string by ignoring spaces in the regular expression: `REGEXP_SUBSTR('oracle', 'o r a c l e', 1, 1, 'x')`
REGEXP_COUNT	Function	Returns the number of times a pattern appears in a string. You specify the string and the pattern. You can also specify the start position and matching options (for example, c for case sensitivity). This function call returns the number of times that e (but not E) appears in the string 'Albert Einstein', starting at character position 7 (that is, one): `REGEXP_COUNT('Albert Einstein', 'e', 7, 'c')`

A string literal in a REGEXP function or condition conforms to the rules of SQL text literals. By default, regular expressions must be enclosed in single quotation marks. If your regular expression includes the single quotation mark, then enter two single quotation marks to represent one single quotation mark within the expression. This technique ensures that the entire expression is interpreted by the SQL function and improves the readability of your code. You can also use the q-quote syntax to define your own character to terminate a text literal. For example, you can delimit your regular expression with the pound sign (#) and then use a single quotation mark within the expression.

Note: If your expression comes from a column or a bind variable, then the preceding rules for quotation marks do not apply.

See Also:

- *Oracle Database SQL Language Reference* for syntax, descriptions, and examples of the `REGEXP` functions and conditions
- *Oracle Database SQL Language Reference* for information about character literals

Oracle Database Support for the POSIX Regular Expression Standard

Oracle Database implementation of regular expressions conforms to these standards:

- IEEE Portable Operating System Interface (POSIX) standard draft 1003.2/D11.2
- Unicode Regular Expression Guidelines of the Unicode Consortium

Oracle Database follows the exact syntax and matching semantics for these operators as defined in the POSIX standard for matching ASCII (English language) data. You can find the POSIX standard draft at this URL:

`http://www.opengroup.org/onlinepubs/007908799/xbd/re.html`

Oracle Database enhances regular expression support in these ways:

- Extends the matching capabilities for multilingual data beyond what is specified in the POSIX standard.
- Adds support for the common PERL regular expression extensions that are not included in the POSIX standard but do not conflict with it. Oracle Database provides built-in support for some heavily used PERL regular expression operators, for example, character class shortcuts, the "nongreedy" modifier, and so on.

Oracle Database supports a set of common metacharacters used in regular expressions. For information about the action of supported metacharacters and related features, see "Metacharacters in Regular Expressions" on page 3-4.

Note: The interpretation of metacharacters differs between tools that support regular expressions. If you are porting regular expressions from another environment to Oracle Database, ensure that the regular expression syntax is supported and the action is what you expect.

Metacharacters in Regular Expressions

Topics:

- POSIX Metacharacters in Oracle Database Regular Expressions
- Multilingual Extensions to POSIX Regular Expression Standard
- PERL-Influenced Extensions to POSIX Regular Expression Standard

POSIX Metacharacters in Oracle Database Regular Expressions

Table 3–2 lists the list of metacharacters supported for use in regular expressions passed to SQL regular expression functions and conditions. These metacharacters conform to the POSIX standard; any differences in action from the standard are noted in the "Description" column.

Table 3–2 POSIX Metacharacters in Oracle Database Regular Expressions

Syntax	Operator Name	Description	Example
`.`	Any Character — Dot	Matches any character in the database character set. If the `n` flag is set, it matches the newline character. The Linux, UNIX, and Windows platforms recognize the newline character as the linefeed character (`\x0a`). The Macintosh platforms recognize the newline character as the carriage return character (`\x0d`). **Note:** In the POSIX standard, this operator matches any English character except NULL and the newline character.	The expression `a.b` matches the strings `abb`, `acb`, and `adb`, but does not match `acc`.
`+`	One or More — Plus Quantifier	Matches one or more occurrences of the preceding subexpression.	The expression `a+` matches the strings `a`, `aa`, and `aaa`, but does not match `bbb`.
`?`	Zero or One — Question Mark Quantifier	Matches zero or one occurrence of the preceding subexpression.	The expression `ab?c` matches the strings `abc` and `ac`, but does not match `abbc`.
`*`	Zero or More — Star Quantifier	Matches zero or more occurrences of the preceding subexpression. By default, a quantifier match is "greedy," because it matches as many occurrences as possible while allowing the rest of the match to succeed.	The expression `ab*c` matches the strings `ac`, `abc`, and `abbc`, but does not match `abb`.
`{`*m*`}`	Interval—Exact Count	Matches exactly *m* occurrences of the preceding subexpression.	The expression `a{3}` matches the strings `aaa`, but does not match `aa`.
`{`*m*`,}`	Interval—At Least Count	Matches at least *m* occurrences of the preceding subexpression.	The expression `a{3,}` matches the strings `aaa` and `aaaa`, but does not match `aa`.
`{`*m*`,`*n*`}`	Interval—Betwee n Count	Matches at least *m*, but not more than *n* occurrences of the preceding subexpression.	The expression `a{3,5}` matches the strings `aaa`, `aaaa`, and `aaaaa`, but does not match `aa`.
`[ ... ]`	Matching Character List	Matches any single character in the list within the brackets. These operators are allowed within the list, but other metacharacters included are treated as literals: ■ Range operator: `-` ■ POSIX character class: `[: :]` ■ POSIX collation element: `[. .]` ■ POSIX character equivalence class: `[= =]` A dash (`-`) is a literal when it occurs first or last in the list, or as an ending range point in a range expression, as in `[#--]`. A right bracket (`]`) is treated as a literal if it occurs first in the list. **Note:** In the POSIX standard, a range includes all collation elements between the start and end of the range in the linguistic definition of the current locale. Thus, ranges are linguistic rather than byte values ranges; the semantics of the range expression are independent of character set. In Oracle Database, the linguistic range is determined by the `NLS_SORT` initialization parameter.	The expression `[abc]` matches the first character in the strings `all`, `bill`, and `cold`, but does not match any characters in `doll`.

Table 3–2 (Cont.) POSIX Metacharacters in Oracle Database Regular Expressions

Syntax	Operator Name	Description	Example
`[^ ... ]`	Nonmatching Character List	Matches any single character not in the list within the brackets. Characters not in the nonmatching character list are returned as a match. See the description of the Matching Character List operator for an account of metacharacters allowed in the character list.	The expression `[^abc]` matches the character `d` in the string `abcdef`, but not the character `a`, `b`, or `c`. The expression `[^abc]+` matches the sequence `def` in the string `abcdef`, but not `a`, `b`, or `c`. The expression `[^a-i]` excludes any character between `a` and `i` from the search result. This expression matches the character `j` in the string `hij`, but does not match any characters in the string `abcdefghi`.
`|`	Or	Matches an alternative.	The expression `a|b` matches character `a` or character `b`.
`( ... )`	Subexpression or Grouping	Treats the expression within parentheses as a unit. The subexpression can be a string of literals or a complex expression containing operators.	The expression `(abc)?def` matches the optional string `abc`, followed by `def`. Thus, the expression matches `abcdefghi` and `def`, but does not match `ghi`.
`\`*n*	Back reference	Matches the nth preceding subexpression, that is, whatever is grouped within parentheses, where *n* is an integer from 1 to 9. The parentheses cause an expression to be remembered; a back reference refers to it. A back reference counts subexpressions from left to right, starting with the opening parenthesis of each preceding subexpression. The expression is invalid if the source string contains fewer than *n* subexpressions preceding the `\`*n*. Oracle Database supports the back reference expression in the regular expression pattern and the replacement string of the `REGEXP_REPLACE` function.	The expression `(abc|def)xy\1` matches the strings `abcxyabc` and `defxydef`, but does not match `abcxydef` or `abcxy`. A backreference enables you to search for a repeated string without knowing the actual string ahead of time. For example, the expression `^(.*)\1$` matches a line consisting of two adjacent instances of the same string.
`\`	Escape Character	Treats the subsequent metacharacter in the expression as a literal. Use a backslash (`\`) to search for a character that is normally treated as a metacharacter. Use consecutive backslashes (`\\`) to match the backslash literal itself.	The expression `\+` searches for the plus character (`+`). It matches the plus character in the string `abc+def`, but does not match `abcdef`.
`^`	Beginning of Line Anchor	Matches the beginning of a string (default). In multiline mode, it matches the beginning of any line within the source string.	The expression `^def` matches `def` in the string `defghi` but does not match `def` in `abcdef`.
`$`	End of Line Anchor	Matches the end of a string (default). In multiline mode, it matches the end of any line within the source string.	The expression `def$` matches `def` in the string `abcdef` but does not match `def` in the string `defghi`.

Table 3–2 (Cont.) POSIX Metacharacters in Oracle Database Regular Expressions

Syntax	Operator Name	Description	Example
`[:class:]`	POSIX Character Class	Matches any character belonging to the specified POSIX character *class*. You can use this operator to search for characters with specific formatting such as uppercase characters, or you can search for special characters such as digits or punctuation characters. The full set of POSIX character classes is supported. **Note:** In English regular expressions, range expressions often indicate a character class. For example, `[a-z]` indicates any lowercase character. This convention is not useful in multilingual environments, where the first and last character of a given character class might not be the same in all languages. Oracle Database supports the character classes in Table 3–3 based on character class definitions in Globalization classification data.	The expression `[[:upper:]]+` searches for one or more consecutive uppercase characters. This expression matches `DEF` in the string `abcDEFghi` but does not match the string `abcdefghi`.
`[.element.]`	POSIX Collating Element Operator	Specifies a collating element to use in the regular expression. The *element* must be a defined collating element in the current locale. Use any collating element defined in the locale, including single-character and multicharacter elements. The `NLS_SORT` initialization parameter determines supported collation elements.This operator lets you use a multicharacter collating element in cases where only one character is otherwise allowed. For example, you can ensure that the collating element `ch`, when defined in a locale such as Traditional Spanish, is treated as one character in operations that depend on the ordering of characters.	The expression `[[.ch.]]` searches for the collating element `ch` and matches `ch` in string `chabc`, but does not match `cdefg`. The expression `[a-[.ch.]]` specifies the range `a` to `ch`.
`[=character=]`	POSIX Character Equivalence Class	Matches all characters that are members of the same character equivalence class in the current locale as the specified *character*. The character equivalence class must occur within a character list, so the character equivalence class is always nested within the brackets for the character list in the regular expression. Usage of character equivalents depends on how canonical rules are defined for your database locale. See *Oracle Database Globalization Support Guide* for more information about linguistic sorting and string searching.	The expression `[[=n=]]` searches for characters equivalent to `n` in a Spanish locale. It matches both `N` and `ñ` in the string `El Niño`.

See Also: *Oracle Database SQL Language Reference* for syntax, descriptions, and examples of the `REGEXP` functions and conditions

Multilingual Extensions to POSIX Regular Expression Standard

When applied to multilingual data, Oracle Database implementation of the POSIX operators extends beyond the matching capabilities specified in the POSIX standard. Table 3–3 shows the relationship of the operators in the POSIX standard.

- The first column lists the supported operators.
- The second column indicates whether the POSIX standard for Basic Regular Expression (BRE) defines the operator.
- The third column indicates whether the POSIX standard for Extended Regular Expression (ERE) defines the operator.
- The fourth column indicates whether the Oracle Database implementation extends the operator's semantics for handling multilingual data.

Oracle Database lets you enter multibyte characters directly, if you have a direct input method, or use functions to compose the multibyte characters. You cannot use the Unicode hexadecimal encoding value of the form `\xxxx`. Oracle Database evaluates the characters based on the byte values used to encode the character, not the graphical representation of the character.

Table 3–3 POSIX and Multilingual Operator Relationships

Operator	POSIX BRE syntax	POSIX ERE Syntax	Multilingual Enhancement
`\`	Yes	Yes	--
`*`	Yes	Yes	--
`+`	--	Yes	--
`?`	--	Yes	--
`|`	--	Yes	--
`^`	Yes	Yes	Yes
`$`	Yes	Yes	Yes
`.`	Yes	Yes	Yes
`[ ]`	Yes	Yes	Yes
`( )`	Yes	Yes	--
`{m}`	Yes	Yes	--
`{m,}`	Yes	Yes	--
`{m,n}`	Yes	Yes	--
`\n`	Yes	Yes	Yes
`[..]`	Yes	Yes	Yes
`[::]`	Yes	Yes	Yes
`[==]`	Yes	Yes	Yes

PERL-Influenced Extensions to POSIX Regular Expression Standard

Table 3–4 describes PERL-influenced metacharacters supported in Oracle Database regular expression functions and conditions. These metacharacters are not in the POSIX standard, but are common at least partly from the popularity of PERL. PERL character class matching is based on the locale model of the operating system, whereas Oracle Database regular expressions are based on the language-specific data of the database. In general, a regular expression involving locale data cannot be expected to produce the same results between PERL and Oracle Database.

Table 3–4 PERL-Influenced Extensions in Oracle Database Regular Expressions

Reg. Exp.	Matches . . .	Example
`\d`	A digit character. It is equivalent to the POSIX class `[[:digit:]]`.	The expression `^\(\d{3}\) \d{3}-\d{4}$` matches `(650) 555-0100` but does not match `650-555-0100`.
`\D`	A nondigit character. It is equivalent to the POSIX class `[^[:digit:]]`.	The expression `\w\d\D` matches `b2b` and `b2_` but does not match `b22`.
`\w`	A word character, which is defined as an alphanumeric or underscore (_) character. It is equivalent to the POSIX class `[[:alnum:]_]`. If you do not want to include the underscore character, you can use the POSIX class `[[:alnum:]]`.	The expression `\w+@\w+(\.\w+)+` matches the string `jdoe@company.co.uk` but not the string `jdoe@company`.
`\W`	A nonword character. It is equivalent to the POSIX class `[^[:alnum:]_]`.	The expression `\w+\W\s\w+` matches the string `to: bill` but not the string `to bill`.
`\s`	A whitespace character. It is equivalent to the POSIX class `[[:space:]]`.	The expression `\(\w\s\w\s\)` matches the string `(a b )` but not the string `(ab)`.
`\S`	A nonwhitespace character. It is equivalent to the POSIX class `[^[:space:]]`.	The expression `\(\w\S\w\S\)` matches the string `(abde)` but not the string `(a b d e)`.
`\A`	Only at the beginning of a string. In multi-line mode, that is, when embedded newline characters in a string are considered the termination of a line, `\A` does not match the beginning of each line.	The expression `\AL` matches only the first `L` character in the string `Line1\nLine2\n`, regardless of whether the search is in single-line or multi-line mode.
`\Z`	Only at the end of a string or before a newline ending a string. In multi-line mode, that is, when embedded newline characters in a string are considered the termination of a line, `\Z` does not match the end of each line.	In the expression `\s\Z`, the `\s` matches the last space in the string `L i n e \n`, regardless of whether the search is in single-line or multi-line mode.
`\z`	Only at the end of a string.	In the expression `\s\z`, the `\s` matches the newline in the string `L i n e \n`, regardless of whether the search is in single-line or multi-line mode.
`*?`	The preceding pattern element 0 or more times ("nongreedy"). This quantifier matches the empty string whenever possible.	The expression `\w*?x\w` is "nongreedy" and so matches `abxc` in the string `abxcxd`. The expression `\w*x\w` is "greedy" and so matches `abxcxd in the string abxcxd`. The expression `\w*?x\w` also matches the string `xa`.
`+?`	The preceding pattern element 1 or more times ("nongreedy").	The expression `\w+?x\w` is "nongreedy" and so matches `abxc` in the string `abxcxd`. The expression `\w+x\w` is "greedy" and so matches `abxcxd in the string abxcxd`. The expression `\w+?x\w` does not match the string `xa`, but does match the string `axa`.
`??`	The preceding pattern element 0 or 1 time ("nongreedy"). This quantifier matches the empty string whenever possible.	The expression `a??aa` is "nongreedy" and matches `aa` in the string `aaaa`. The expression `a?aa` is "greedy" and so matches `aaa` in the string `aaaa`.

Table 3–4 (Cont.) PERL-Influenced Extensions in Oracle Database Regular Expressions

Reg. Exp.	Matches . . .	Example
`{n}?`	The preceding pattern element exactly `n` times ("nongreedy"). In this case `{n}?` is equivalent to `{n}`.	The expression `(a|aa){2}?` matches `aa` in the string `aaaa`.
`{n,}?`	The preceding pattern element at least `n` times ("nongreedy").	The expression `a{2,}?` is "nongreedy" and matches `aa` in the string `aaaaa`. The expression `a{2,}` is "greedy" and so matches `aaaaa`.
`{n,m}?`	At least `n` but not more than `m` times ("nongreedy"). `{0,m}?` matches the empty string whenever possible.	The expression `a{2,4}?` is "nongreedy" and matches `aa` in the string `aaaaa`. The expression `a{2,4}` is "greedy" and so matches `aaaa`.

Oracle Database regular expression functions and conditions support the pattern matching modifiers described in Table 3–5.

Table 3–5 Pattern Matching Modifiers

Mod.	Description	Example
`i`	Specifies case-insensitive matching.	This regular expression returns `AbCd`: `REGEXP_SUBSTR('AbCd', 'abcd', 1, 1, 'i')`
`c`	Specifies case-sensitive matching.	This regular expression fails to match: `REGEXP_SUBSTR('AbCd', 'abcd', 1, 1, 'c')`
`n`	Allows the period (.), which by default does not match newlines, to match the newline character.	This regular expression matches the string only because the `n` flag is specified: `REGEXP_SUBSTR('a'||CHR(10)||'d', 'a.d', 1, 1, 'n')`
`m`	Performs the search in multi-line mode. The metacharacter `^` and `$` signify the start and end, respectively, of any line anywhere in the source string, rather than only at the start or end of the entire source string.	This regular expression returns `ac`: `REGEXP_SUBSTR('ab'||CHR(10)||'ac', '^a.', 1, 2, 'm')`
`x`	Ignores whitespace characters in the regular expression. By default, whitespace characters match themselves.	This regular expression returns `abcd`: `REGEXP_SUBSTR('abcd', 'a b c d', 1, 1, 'x')`

Using Regular Expressions in SQL Statements: Scenarios

Scenarios:

- Using a Constraint to Enforce a Phone Number Format
- Using Back References to Reposition Characters

Using a Constraint to Enforce a Phone Number Format

Regular expressions are useful for enforcing constraints—for example, to ensure that phone numbers are entered into the database in a standard format. Example 3–1 creates a `contacts` table and adds a `CHECK` constraint to the `p_number` column to enforce this format mask:

```
(XXX) XXX-XXXX
```

Example 3–1 Enforcing a Phone Number Format with Regular Expressions

```
DROP TABLE contacts;
CREATE TABLE contacts (
```

```
  l_name    VARCHAR2(30),
  p_number  VARCHAR2(30)
  CONSTRAINT c_contacts_pnf
  CHECK (REGEXP_LIKE (p_number, '^\(\d{3}\) \d{3}-\d{4}$'))
);
```

Table 3–6 explains the elements of the regular expression.

Table 3–6 Explanation of the Regular Expression Elements in Example 3–1

Regular Expression Element	Matches . . .
^	The beginning of the string.
\(	A left parenthesis. The backward slash (\) is an escape character that indicates that the left parenthesis after it is a literal rather than a grouping expression.
\d{3}	Exactly three digits.
\)	A right parenthesis. The backward slash (\) is an escape character that indicates that the right parenthesis after it is a literal rather than a grouping expression.
(space character)	A space character.
\d{3}	Exactly three digits.
-	A hyphen.
\d{4}	Exactly four digits.
$	The end of the string.

Example 3–2 Inserting Phone Numbers in Correct and Incorrect Formats

These are correct:

```
INSERT INTO contacts (p_number) VALUES('(650) 555-0100');
INSERT INTO contacts (p_number) VALUES('(215) 555-0100');
```

These generate CHECK constraint errors:

```
INSERT INTO contacts (p_number) VALUES('650 555-0100');
INSERT INTO contacts (p_number) VALUES('650 555 0100');
INSERT INTO contacts (p_number) VALUES('650-555-0100');
INSERT INTO contacts (p_number) VALUES('(650)555-0100');
INSERT INTO contacts (p_number) VALUES(' (650) 555-0100');
```

Using Back References to Reposition Characters

As explained in Table 3–2, back references store matched subexpressions in a temporary buffer, enabling you to reposition characters. You access buffers with the *n* notation, where *n* is a number in the range from 1 through 9. Each subexpression is enclosed in parentheses, and its characters are numbered from left to right.

Example 3–3 creates a table, populates it with names in different formats, and uses a query that repositions names that are in the format "first middle last" to the format "last, first middle". It ignores names not in the format "first middle last". The elements of the regular expression in the query are explained in Table 3–7.

Example 3–3 Using Back References to Reposition Characters

Create and populate table:

```
DROP TABLE famous_people;
CREATE TABLE famous_people (names VARCHAR2(20));
INSERT INTO famous_people (names) VALUES ('John Quincy Adams');
INSERT INTO famous_people (names) VALUES ('Harry S. Truman');
INSERT INTO famous_people (names) VALUES ('John Adams');
INSERT INTO famous_people (names) VALUES (' John Quincy Adams');
INSERT INTO famous_people (names) VALUES ('John_Quincy_Adams');
```

SQL*Plus formatting command:

```
COLUMN "names after regexp" FORMAT A20
```

Repositioning query:

```
SELECT names "names",
  REGEXP_REPLACE(names, '^(\S+)\s(\S+)\s(\S+)$', '\3, \1 \2')
    AS "names after regexp"
FROM famous_people
ORDER BY "names";
```

Result:

```
names                names after regexp
-------------------- --------------------
 John Quincy Adams    John Quincy Adams
Harry S. Truman      Truman, Harry S.
John Adams           John Adams
John Quincy Adams    Adams, John Quincy
John_Quincy_Adams    John_Quincy_Adams

5 rows selected.
```

Table 3–7 Explanation of the Regular Expression Elements in Example 3–3

Regular Expression Element	Description
^	Matches the beginning of the string.
$	Matches the end of the string.
(\S+)	Matches one or more nonspace characters. The parentheses are not escaped so they function as a grouping expression.
\s	Matches a whitespace character.
\1	Substitutes the first subexpression, that is, the first group of parentheses in the matching pattern.
\2	Substitutes the second subexpression, that is, the second group of parentheses in the matching pattern.
\3	Substitutes the third subexpression, that is, the third group of parentheses in the matching pattern.
,	Inserts a comma character.

4

Using Indexes in Database Applications

This chapter explains how to use indexes in database applications.

Topics:

- Privileges Needed to Create Indexes
- Guidelines for Application-Specific Indexes
- Examples of Creating Basic Indexes
- When to Use Domain Indexes
- When to Use Function-Based Indexes

See Also:

- *Oracle Database Administrator's Guide* for information about creating and managing indexes
- *Oracle Database Performance Tuning Guide* for detailed information about using indexes
- *Oracle Database SQL Language Reference* for the syntax of statements to work with indexes
- *Oracle Database Administrator's Guide* for information about creating hash clusters to improve performance, as an alternative to indexing

Privileges Needed to Create Indexes

When using indexes in an application, you might need to ask the DBA to grant privileges or make changes to initialization parameters.

To create an index, you must own, or have the `INDEX` object privilege for, the corresponding table. The schema that contains the index must also have a quota for the tablespace intended to contain the index, or the `UNLIMITED TABLESPACE` system privilege. To create an index in another user's schema, you must have the `CREATE ANY INDEX` system privilege.

Guidelines for Application-Specific Indexes

You can create indexes on columns to speed up queries. Indexes provide faster access to data for operations that return a small portion of a table's rows.

In general, create an index on a column in any of these situations:

- The column is queried frequently.
- A referential constraint exists on the column.
- A `UNIQUE` key constraint exists on the column.

You can create an index on any column; however, if the column is not used in any of these situations, creating an index on the column does not increase performance and the index takes up resources unnecessarily.

Although the database creates an index for you on a column with a constraint, explicitly creating an index on such a column is recommended.

You can use these techniques to determine which columns are best candidates for indexing:

- Use the `EXPLAIN PLAN` feature to show a theoretical execution plan of a given query statement.
- Use the dynamic performance view `V$SQL_PLAN` to determine the actual execution plan used for a given query statement.

Sometimes, if an index is not being used by default and it would be more efficient to use that index, you can use a query hint so that the index is used.

See Also:

- *Oracle Database Performance Tuning Guide* for information about using the `V$SQL_PLAN` view, the `EXPLAIN PLAN` statement, query hints, and measuring the performance benefits of indexes
- *Oracle Database Reference* for general information about the `V$SQL_PLAN` view

Topics:

- Which Come First, Data or Indexes?
- Create a Temporary Table Space Before Creating Indexes
- Index the Correct Tables and Columns
- Limit the Number of Indexes for Each Table
- Choose Column Order in Composite Indexes
- Gather Index Statistics
- Drop Unused Indexes

Which Come First, Data or Indexes?

Typically, you insert or load data into a table (using SQL*Loader or Import) before creating indexes. Otherwise, the overhead of updating the index slows down the insert or load operation. The exception to this rule is that you must create an index for a cluster before you insert any data into the cluster.

Create a Temporary Table Space Before Creating Indexes

When you create an index on a table that has data, Oracle Database must use sort space to create the index. The database uses the sort space in memory allocated for the creator of the index (the amount for each user is determined by the initialization parameter `SORT_AREA_SIZE`), but the database must also swap sort information to

and from temporary segments allocated on behalf of the index creation. If the index is extremely large, Oracle recommends following these steps:

1. Create a temporary tablespace using the `CREATE TABLESPACE` statement.
2. Use the `TEMPORARY TABLESPACE` option of the `ALTER USER` statement to make this your temporary tablespace.
3. Create the index using the `CREATE INDEX` statement.
4. Drop this tablespace using the `DROP TABLESPACE` statement. Then use the `ALTER USER` statement to reset your temporary tablespace to your original temporary tablespace.

Under certain conditions, you can load data into a table with the SQL*Loader "direct path load", and an index can be created as data is loaded.

> **See Also:** *Oracle Database Utilities* for information about direct path load

Index the Correct Tables and Columns

Use these guidelines for determining when to create an index:

- Create an index if you frequently want to retrieve less than about 15% of the rows in a large table. This threshold percentage varies greatly, however, according to the relative speed of a table scan and how clustered the row data is about the index key. The faster the table scan, the lower the percentage; the more clustered the row data, the higher the percentage.
- Index columns that are used for joins to improve join performance.
- Primary and unique keys automatically have indexes, but you might want to create an index on a foreign key; see Chapter 5, "Maintaining Data Integrity in Database Applications," for more information.
- Small tables do not require indexes; if a query is taking too long, then the table might have grown from small to large.

Some columns are strong candidates for indexing. Columns with one or more of these characteristics are good candidates for indexing:

- Values are unique in the column, or there are few duplicates.
- There is a wide range of values (good for regular indexes).
- There is a small range of values (good for bitmap indexes).
- The column contains many nulls, but queries often select all rows having a value. In this case, a comparison that matches all the non-null values, such as:

  ```
  WHERE COL_X >= -9.99 *power(10,125)
  ```

  is preferable to

  ```
  WHERE COL_X IS NOT NULL
  ```

  This is because the first uses an index on `COL_X` (if `COL_X` is a numeric column).

Columns with these characteristics are less suitable for indexing:

- There are many nulls in the column and you do not search on the non-null values.

`LONG` and `LONG RAW` columns cannot be indexed.

The size of a single index entry cannot exceed roughly one-half (minus some overhead) of the available space in the data block. Consult with the database administrator for assistance in determining the space required by an index.

Limit the Number of Indexes for Each Table

The more indexes, the more overhead is incurred as the table is altered. When rows are inserted or deleted, all indexes on the table must be updated. When a column is updated, all indexes on the column must be updated.

You must weigh the performance benefit of indexes for queries against the performance overhead of updates. For example, if a table is primarily read-only, you might use more indexes; but, if a table is heavily updated, you might use fewer indexes.

Choose Column Order in Composite Indexes

Although you can specify columns in any order in the `CREATE INDEX` statement, the order of columns in the `CREATE INDEX` statement can affect query performance. In general, put the column expected to be used most often first in the index. You can create a composite index (using several columns), and use the same index for queries that reference all or some of these columns.

Example 4–1 VENDOR_PARTS Table

```
DROP TABLE vendor_parts;
CREATE TABLE vendor_parts (
  vendor_id  VARCHAR2(9),
  part_no    VARCHAR2(7),
  unit_cost  REAL
);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1012, 10440, .25);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1012, 10441, .39);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1012, 457, 4.95);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1010, 10440, .27);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1010, 457, 5.12);

INSERT INTO vendor_parts  (vendor_id, part_no, unit_cost)
VALUES (1220, 8300, 1.33);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1012, 8300, 1.19);

INSERT INTO vendor_parts (vendor_id, part_no, unit_cost)
VALUES (1292, 457, 5.28);
```

Query:

```
SELECT * FROM vendor_parts
```

```
ORDER BY vendor_id;
```

Result:

```
VENDOR_ID PART_NO  UNIT_COST
--------- ------- ----------
1010      10440          .27
1010      457           5.12
1012      457           4.95
1012      8300          1.19
1012      10441          .39
1012      10440          .25
1220      8300          1.33
1292      457           5.28

8 rows selected.
```

Assume that there are five vendors, and each vendor has about 1000 parts.

Suppose that the `VENDOR_PARTS` table is commonly queried by SQL statements such as:

```
SELECT * FROM vendor_parts
WHERE part_no = 457 AND vendor_id = 1012
ORDER BY vendor_id;
```

Result:

```
VENDOR_ID PART_NO  UNIT_COST
--------- ------- ----------
1012      457           4.95

1 row selected.
```

To increase the performance of such queries, you might create a composite index putting the most selective column first; that is, the column with the most values:

```
CREATE INDEX ind_vendor_id
ON vendor_parts (part_no, vendor_id);
```

Composite indexes speed up queries that use the leading portion of the index. So in this example, the performance of queries with `WHERE` clauses using only the `PART_NO` column improve also. Because there are only five distinct values, placing a separate index on `VENDOR_ID` serves no purpose.

Gather Index Statistics

The database can use indexes more effectively when it has statistical information about the tables involved in the queries. You or the DBA can periodically gather statistics by invoking procedures such as `DBMS_STATS.GATHER_TABLE_STATISTICS` and `DBMS_STATS.GATHER_SCHEMA_STATISTICS`. For information about these procedures, see *Oracle Database PL/SQL Packages and Types Reference*.

Drop Unused Indexes

You might drop an index if:

- It does not speed up queries. The table might be very small, or there might be many rows in the table but very few index entries.
- The queries in your applications do not use the index.

To find out if an index is being used, you can monitor it. If you see that the index is never used, rarely used, or used in a way that seems to provide no benefit, you can either drop it immediately or you can make it invisible until you are sure that you do not need it, and then drop it. If you discover that you do need the invisible index, you can make it visible again.

When you drop an index, all extents of the index's segment are returned to the containing tablespace and become available for other objects in the tablespace.

To drop an index, use the SQL statement `DROP INDEX`. For example, this statement drops the index named `Emp_name`:

```
DROP INDEX Emp_ename;
```

If you drop a table, then all associated indexes are dropped.

To drop an index, the index must be contained in your schema or you must have the `DROP ANY INDEX` system privilege.

See Also:

- *Oracle Database Administrator's Guide* for information about monitoring index usage
- *Oracle Database Administrator's Guide* for information about making indexes invisible
- *Oracle Database SQL Language Reference* for information about the `DROP INDEX` statement

Examples of Creating Basic Indexes

You can create an index for a table to improve the performance of queries issued against the corresponding table. You can also create an index for a cluster. You can create a *composite* index on multiple columns up to a maximum of 32 columns. A composite index key cannot exceed roughly one-half (minus some overhead) of the available space in the data block.

Oracle Database automatically creates an index to enforce a `UNIQUE` or `PRIMARY KEY` constraint. In general, it is better to create such constraints to enforce uniqueness, instead of using the obsolete `CREATE UNIQUE INDEX` syntax.

Use the SQL statement `CREATE INDEX` to create an index.

Example 4–2 Creating Indexes

Create index for single column, to speed up queries that test that column:

```
CREATE INDEX emp_phone ON EMPLOYEES(PHONE_NUMBER);
```

Create index for single column, specifying some physical attributes for index:

```
CREATE INDEX emp_lastname ON EMPLOYEES (LAST_NAME)
  STORAGE (
    INITIAL     20K
    NEXT        20k
    PCTINCREASE 75
  )
  PCTFREE  0;
```

Create index for two columns, to speed up queries that test either first column or both columns:

```
CREATE INDEX emp_id_email ON EMPLOYEES(EMPLOYEE_ID, EMAIL);
```

For query that sorts on `UPPER(LASTNAME)`, index on `LAST_NAME` column does not speed up operation, and might be slow to invoke function for each result row. Create function-based index that precomputes the result of the function for each column value,speeding up queries that use the function for searching or sorting:

```
CREATE INDEX emp_upper_lastname ON EMPLOYEES(UPPER(LAST_NAME));
```

When to Use Domain Indexes

Domain indexes are appropriate for special-purpose applications implemented using data cartridges. The domain index helps to manipulate complex data, such as spatial, audio, or video data. If you must develop such an application, see *Oracle Database Data Cartridge Developer's Guide*.

Oracle Database supplies specialized data cartridges to help manage these kinds of complex data. So, if you must create a search engine, or a geographic information system, you can do much of the work simply by creating the right kind of index.

When to Use Function-Based Indexes

A function-based index is an index built on an expression. It extends your indexing capabilities beyond indexing on a column. A function-based index increases the variety of ways in which you can access data.

Note:

- The index is more effective if you gather statistics for the table or schema, using the procedures in the `DBMS_STATS` package.
- The index cannot contain any null values. Either ensure that the appropriate columns contain no null values, or use the `NVL` function in the index expression to substitute some other value for nulls.

The expression indexed by a function-based index can be an arithmetic expression or an expression that contains a PL/SQL function, package function, C callout, or SQL function. Function-based indexes also support linguistic sorts based on collation keys, efficient linguistic collation of SQL statements, and case-insensitive sorts.

Like other indexes, function-based indexes improve query performance. For example, if you must access a computationally complex expression often, then you can store it in an index. Then when you must access the expression, it is available. You can find a detailed description of the advantages of function-based indexes in "Advantages of Function-Based Indexes" on page 4-8.

Function-based indexes have all of the same properties as indexes on columns. Unlike indexes on columns that can be used by both cost-based and rule-based optimization, however, function-based indexes can be used by only by cost-based optimization. Other restrictions on function-based indexes are described in "Restrictions on Function-Based Indexes" on page 4-10.

See Also:

- *Oracle Database Concepts* for general information about function-based indexes
- *Oracle Database Administrator's Guide* for information about creating function-based indexes

Topics:

- Advantages of Function-Based Indexes
- Restrictions on Function-Based Indexes
- Examples of Function-Based Indexes

Advantages of Function-Based Indexes

Function-based indexes:

- Increase the number of situations where the optimizer can perform a range scan instead of a full table scan (as in Example 4–3).

  Range scans typically produce fast response times if the predicate selects less than 15% of the rows of a large table. The optimizer can estimate how many rows are selected by expressions more accurately if the expressions are materialized in a function-based index. (Expressions of function-based indexes are represented as virtual columns and `ANALYZE` can build histograms on such columns.)

- Precompute the value of a computationally intensive function and store it in the index.

  An index can store computationally intensive expression that you access often. When you must access a value, it is available, greatly improving query execution performance.

- Create indexes on object columns and `REF` columns.

  Methods that describe objects can be used as functions on which to build indexes. For example, you can use the `MAP` method to build indexes on an ADT column.

- Create more powerful sorts.

  You can perform case-insensitive sorts with the `UPPER` and `LOWER` functions, descending order sorts with the `DESC` keyword, and linguistic-based sorts with the `NLSSORT` function.

Note: Oracle Database sorts columns with the `DESC` keyword in descending order. Such indexes are treated as function-based indexes. Descending indexes cannot be bitmapped or reverse, and cannot be used in bitmapped optimizations. To get the `DESC` functionality before Oracle Database version 8, remove the `DESC` keyword from the `CREATE INDEX` statement.

In Example 4–3, an index is built on (`Column_a` + `Column_b`); therefore, the expression in the `WHERE` clause of the `SELECT` statement allows the optimizer to perform a range scan instead of a full table scan.

Example 4–3 Function-Based Index Allows Optimizer to Perform Range Scan

```
DROP TABLE Example_tab;
```

```
CREATE TABLE Example_tab (
  Column_a INTEGER,
  Column_b INTEGER
);

INSERT INTO Example_tab (Column_a, Column_b)
VALUES (1, 2);

INSERT INTO Example_tab (Column_a, Column_b)
VALUES (2, 4);

INSERT INTO Example_tab (Column_a, Column_b)
VALUES (3, 6);

INSERT INTO Example_tab (Column_a, Column_b)
VALUES (4, 8);

INSERT INTO Example_tab (Column_a, Column_b)
VALUES (5, 10);
```

Query:

```
SELECT * FROM Example_tab ORDER BY Column_a;
```

Result:

```
  COLUMN_A   COLUMN_B
---------- ----------
         1          2
         2          4
         3          6
         4          8
         5         10

5 rows selected.
```

Create index:

```
CREATE INDEX Idx ON Example_tab(Column_a + Column_b);
```

Query:

```
SELECT * FROM Example_tab
WHERE Column_a + Column_b < 10
ORDER BY Column_a;
```

Result:

```
  COLUMN_A   COLUMN_B
---------- ----------
         1          2
         2          4
         3          6

3 rows selected.
```

In Example 4–4:

- The function-based index `Distance_index` calls the object method `Distance_from_equator` for each city in a table. The method is applied to the object column `Reg_Obj`. A query uses `Distance_index` to quickly find cities that are

more than 1000 miles from the equator. (The table is not populated for the example, so the query returns no rows.)

- The function-based index `Compare_index` stores the temperature delta and the maximum temperature. The result of the delta is sorted in descending order. A query uses `Compare_index` to quickly find table rows where the temperature delta is less than 20 and the maximum temperature is greater than 75. (The table is not populated for the example, so the query returns no rows.)

Example 4–4 Function-Based Indexes

```
DROP TABLE Weatherdata_tab;
CREATE TABLE Weatherdata_tab (
  Reg_obj INTEGER,
  Maxtemp INTEGER,
  Mintemp INTEGER
);

CREATE OR REPLACE FUNCTION Distance_from_equator (
  Reg_obj IN INTEGER
) RETURN INTEGER
  DETERMINISTIC
IS
BEGIN
  RETURN(3000);
END;
/
```

Create index:

```
CREATE INDEX Distance_index
ON Weatherdata_tab (Distance_from_equator (Reg_obj));
```

Query:

```
SELECT * FROM Weatherdata_tab
WHERE (Distance_from_equator (Reg_Obj)) > '1000';
```

Result:

```
no rows selected
```

Create index:

```
CREATE INDEX Compare_index
  2    ON Weatherdata_tab ((Maxtemp - Mintemp) DESC, Maxtemp);
```

Query:

```
SELECT * FROM Weatherdata_tab
WHERE ((Maxtemp - Mintemp) < '20' AND Maxtemp > '75');
```

Result:

```
no rows selected
```

Restrictions on Function-Based Indexes

Function-based indexes have these restrictions:

- Only cost-based optimization can use function-based indexes. Remember to invoke DBMS_STATS.GATHER_TABLE_STATISTICS or DBMS_STATS.GATHER_SCHEMA_STATISTICS, for the function-based index to be effective.
- Any top-level or package-level PL/SQL functions that are used in the index expression must be declared as DETERMINISTIC. That is, they always return the same result given the same input, for example, the UPPER function. You must ensure that the subprogram really is deterministic, because Oracle Database does not check that the assertion is true.

  These semantic rules demonstrate how to use the keyword DETERMINISTIC:

  - You can declare a top level subprogram as DETERMINISTIC.
  - You can declare a PACKAGE level subprogram as DETERMINISTIC in the PACKAGE specification but not in the PACKAGE BODY. An exception is raised if DETERMINISTIC is used inside a PACKAGE BODY.
  - You can declare a private subprogram (declared inside another subprogram or a PACKAGE BODY) as DETERMINISTIC.
  - A DETERMINISTIC subprogram can invoke another subprogram whether the invoked subprogram is declared as DETERMINISTIC or not.
- If you change the semantics of a DETERMINISTIC function and recompile it, then existing function-based indexes and materialized views report results for the prior version of the function. Thus, if you change the semantics of a function, you must manually rebuild any dependent function-based indexes and materialized views.
- Expressions in a function-based index cannot contain any aggregate functions. The expressions must reference only columns in a row in the table.
- You must analyze the table or index before the index is used.
- Bitmap optimizations cannot use descending indexes.
- Function-based indexes are not used when OR-expansion is done.
- The index function cannot be marked NOT NULL. To avoid a full table scan, you must ensure that the query cannot fetch null values.
- Function-based indexes cannot use expressions that return VARCHAR2 or RAW data types of unknown length from PL/SQL functions. A workaround is to limit the size of the function's output by indexing a substring of known length. For example:

```
CREATE OR REPLACE FUNCTION initials (
  name  IN VARCHAR2
) RETURN VARCHAR2
  DETERMINISTIC
IS
BEGIN
  RETURN('A. J.');
END;
/

/* Invoke SUBSTR both when creating index and when referencing
   function in queries. */

CREATE INDEX func_substr_index ON
EMPLOYEES(SUBSTR(initials(FIRST_NAME),1,10));

SELECT SUBSTR(initials(FIRST_NAME),1,10) FROM EMPLOYEES;
```

See Also: *Oracle Database PL/SQL Language Reference* for CREATE FUNCTION restrictions

Examples of Function-Based Indexes

- Function-Based Index for Case-Insensitive Searches
- Precomputing Arithmetic Expressions with a Function-Based Index
- Function-Based Index for Language-Dependent Sorting

Function-Based Index for Case-Insensitive Searches

This statement allows faster case-insensitive searches in table EMP_TAB.

```
CREATE INDEX emp_lastname ON EMPLOYEES (UPPER(LAST_NAME));
```

The SELECT statement uses the function-based index on UPPER(LAST_NAME) to return all of the employees with name like :KEYCOL.

```
SELECT first_name, last_name
FROM EMPLOYEES
WHERE UPPER(LAST_NAME) LIKE 'J%S_N';
```

Result:

```
FIRST_NAME           LAST_NAME
-------------------- -------------------------
Charles              Johnson

1 row selected.
```

Precomputing Arithmetic Expressions with a Function-Based Index

This statement computes a value for each row using columns A, B, and C, and stores the results in the index.

```
DROP TABLE Fbi_tab;
CREATE TABLE Fbi_tab (
  a INTEGER,
  b INTEGER,
  c INTEGER
);

CREATE INDEX Idx ON Fbi_tab (a + b * (c - 1), a, b);
```

The SELECT statement can either use index range scan (because the expression is a prefix of index Idx) or index fast full scan (which might be preferable if the index has specified a high parallel degree).

```
SELECT a FROM Fbi_tab WHERE a + b * (c - 1) < 100;
```

Function-Based Index for Language-Dependent Sorting

This example uses a function-based index to sort based on the collation order for a national language. The NLSSORT function returns a sort key for each name, using the collation sequence GERMAN.

```
DROP TABLE nls_tab;
CREATE TABLE nls_tab (NAME VARCHAR2(80));

CREATE INDEX nls_index
  ON nls_tab (NLSSORT(NAME, 'NLS_SORT = GERMAN'));
```

The SELECT statement selects all of the contents of the table and orders it by NAME. The rows are ordered using the German collation sequence. The Globalization Support parameters are not needed in the SELECT statement, because in a German session, NLS_SORT is set to German and NLS_COMP is set to ANSI.

```
SELECT * FROM nls_tab
WHERE NAME IS NOT NULL
ORDER BY NAME;
```

5

Maintaining Data Integrity in Database Applications

This chapter explains how to use constraints to enforce the business rules associated with your database and prevent the entry of invalid information into tables.

Topics:

- Overview of Constraints
- Enforcing Referential Integrity with Constraints
- Minimizing Space and Time Overhead for Indexes Associated with Constraints
- Guidelines for Indexing Foreign Keys
- Referential Integrity in a Distributed Database
- When to Use CHECK Constraints
- Examples of Defining Constraints
- Enabling and Disabling Constraints
- Modifying Constraints
- Renaming Constraints
- Dropping Constraints
- Managing FOREIGN KEY Constraints
- Viewing Information About Constraints

Overview of Constraints

You can define constraints to enforce business rules on data in your tables. Business rules specify conditions and relationships that must always be true, or must always be false. Because each company defines its own policies about things like salaries, employee numbers, inventory tracking, and so on, you can specify a different set of rules for each database table.

When an integrity constraint applies to a table, all data in the table must conform to the corresponding rule. When you issue a SQL statement that modifies data in the table, Oracle Database ensures that the new data satisfies the integrity constraint, without any checking within your program.

Enforcing Business Rules with Constraints

You can enforce rules by defining constraints more reliably than by adding logic to your application. Oracle Database can check that all the data in a table obeys an integrity constraint faster than an application can.

For example, to ensure that each employee works for a valid department:

1. Create tables `dept_tab` and `emp_tab`:

```
DROP TABLE dept_tab;
CREATE TABLE dept_tab (
  deptname VARCHAR2(20),
  deptno INTEGER
);

DROP TABLE emp_tab;
CREATE TABLE emp_tab (
  empname VARCHAR2(80),
  empno INTEGER, deptno INTEGER
);
```

2. Create a rule that all values in the department table are unique:

```
ALTER TABLE dept_tab
  ADD PRIMARY KEY (deptno);
```

3. Create a rule that every department listed in the employee table must match a value in the department table:

```
ALTER TABLE emp_tab
ADD FOREIGN KEY (deptno)
REFERENCES dept_tab(deptno);
```

When you add an employee record to the table, Oracle Database automatically checks that its department number appears in the department table.

To enforce this rule without constraints, you can use a trigger to query the department table and test that each employee's department is valid. This method is less reliable than using constraints, because `SELECT` in Oracle Database uses consistent read (CR), so the query might miss uncommitted changes from other transactions.

Enforcing Business Rules with Application Logic

You might enforce business rules through both application logic and constraints, if you can filter out bad data before attempting an insert or update. This might let you provide instant feedback to the user, and reduce the load on the database. This technique is appropriate when you can determine that data values are wrong or out of range without checking against any data in the table.

Creating Indexes for Use with Constraints

All enabled unique and primary keys require corresponding indexes. Create these indexes by hand, rather than letting the database create them. Note that:

- Constraints use existing indexes where possible, rather than creating indexes.
- Unique and primary keys can use both nonunique and unique indexes. They can even use only the first few columns of nonunique indexes.
- At most one unique or primary key can use each nonunique index.

- The column orders in the index and the constraint need not match.
- If you must check whether an index is used by a constraint, for example when you want to drop the index, the object number of the index used by a unique or primary key constraint is stored in CDEF$.ENABLED for that constraint. It is not shown in any static data dictionary view or dynamic performance view.
- Oracle Database does not automatically index foreign keys.

When to Use NOT NULL Constraints

By default, all columns can contain null values. Define NOT NULL constraints only for columns that always require values. For example, an employee's manager or hire date might be temporarily omitted. Some employees might not have a commission. Columns like these must not have NOT NULL constraints. However, an employee name might be required from the very beginning, and you can enforce this rule with a NOT NULL integrity constraint.

NOT NULL constraints are often combined with other constraints to further restrict the values that can exist in specific columns. For example, the combination of NOT NULL and UNIQUE constraints forces the input of values in the UNIQUE key, eliminating the possibility that a new row's data conflicts with an existing row's data.

Because Oracle Database indexes do not store keys that are all null, to allow index-only scans of the table or some other operation that requires indexing all rows, you must put a NOT NULL constraint on at least one indexed column.

> **See Also:** "Defining Relationships Between Parent and Child Tables" on page 5-9

Specify a NOT NULL constraint like this:

```
ALTER TABLE table_name MODIFY column_name NOT NULL;
```

Example 5–1 uses the SQL*Plus command DESCRIBE to show which columns of the DEPARTMENTS table have NOT NULL constraints, and then shows what happens if you try to insert NULL values in columns that have NOT NULL constraints.

Example 5–1 Inserting NULL Values into Columns with NOT NULL Constraints

```
DESCRIBE DEPARTMENTS;
```

Result:

```
 Name                                      Null?    Type
 ----------------------------------------- -------- ------------

 DEPARTMENT_ID                             NOT NULL NUMBER(4)
 DEPARTMENT_NAME                           NOT NULL VARCHAR2(30)
 MANAGER_ID                                         NUMBER(6)
 LOCATION_ID                                        NUMBER(4)
```

Try to insert NULL into DEPARTMENT_ID column:

```
INSERT INTO DEPARTMENTS (
 DEPARTMENT_ID, DEPARTMENT_NAME, MANAGER_ID, LOCATION_ID
)
VALUES (NULL, 'Sales', 200, 1700);
```

Result:

```
VALUES (NULL, 'Sales', 200, 1700)
        *
ERROR at line 4:
ORA-01400: cannot insert NULL into ("HR"."DEPARTMENTS"."DEPARTMENT_ID")
```

Omitting a value for a column that cannot be NULL is the same as assigning it the value NULL:

```
INSERT INTO DEPARTMENTS (
  DEPARTMENT_NAME, MANAGER_ID, LOCATION_ID
)
VALUES ('Sales', 200, 1700);
```

Result:

```
INSERT INTO DEPARTMENTS (
*
ERROR at line 1:
ORA-01400: cannot insert NULL into ("HR"."DEPARTMENTS"."DEPARTMENT_ID")
```

When to Use Default Column Values

Assign default values to columns that contain typical values. For example, in the DEPT_TAB table, if most departments are located in New York, then the default value for the LOC column can be set to NEW YORK.

Default values can help avoid errors where there is a number, such as zero, that applies to a column that has no entry. For example, a default value of zero can simplify testing, by changing a test like this:

```
IF sal IS NOT NULL AND sal < 50000
```

to the simpler form:

```
IF sal < 50000
```

Depending upon your business rules, you might use default values to represent zero or false, or leave the default values as NULL to signify an unknown value.

Defaults are also useful when you use a view to make a subset of a table's columns visible. For example, you might allow users to insert rows through a view. The base table might also have a column named inserter, not included in the definition of the view, to log the user that inserts each row. To record the user name automatically, define a default value that invokes the USER function:

```
CREATE TABLE audit_trail (
  value1   NUMBER,
  value2   VARCHAR2(32),
  inserter VARCHAR2(30) DEFAULT USER);
```

Setting Default Column Values

Default values can be defined using any literal, or almost any expression, including calls to these functions:

- SYSDATE
- SYS_CONTEXT
- USER
- USERENV

- UID

Default values cannot include expressions that refer to a sequence, PL/SQL function, column, LEVEL, ROWNUM, or PRIOR. The data type of a default literal or expression must match or be convertible to the column data type.

Sometimes the default value is the result of a SQL function. For example, a call to SYS_CONTEXT can set a different default value depending on conditions such as the user name. To be used as a default value, a SQL function must have parameters that are all literals, cannot reference any columns, and cannot invoke any other functions.

If you do not explicitly define a default value for a column, the default for the column is implicitly set to NULL.

You can use the keyword DEFAULT within an INSERT statement instead of a literal value, and the corresponding default value is inserted.

Choosing a Primary Key for a Table

Each table can have one primary key, which uniquely identifies each row in a table and ensures that no duplicate rows exist. When selecting a primary key, use these guidelines:

- Whenever practical, use a column containing a sequence number. This satisfies all the other guidelines.
- Choose a column whose data values are unique, because the purpose of a primary key is to uniquely identify each row of the table.
- Choose a column whose data values never change. A primary key value is only used to identify a row in the table, and its data must never be used for any other purpose.
- Choose a column that does not contain any nulls. A PRIMARY KEY constraint, by definition, does not allow any row to contain a null in any column that is part of the primary key.
- Choose a column that is short and numeric. Short primary keys are easy to type. You can use sequence numbers to easily generate numeric primary keys.
- Minimize your use of composite primary keys. Although composite primary keys are allowed, they do not satisfy all of the other recommendations. For example, composite primary key values are long and cannot be assigned by sequence numbers.

When to Use UNIQUE Constraints

Choose columns for unique keys carefully. The purpose of these constraints is different from that of primary keys. Unique key constraints are appropriate for any column where duplicate values are not allowed. Primary keys identify each row of the table uniquely, and typically contain values that have no significance other than being unique. Figure 5–1 shows an example of a table with a unique key constraint.

Figure 5–1 Table with a UNIQUE Constraint

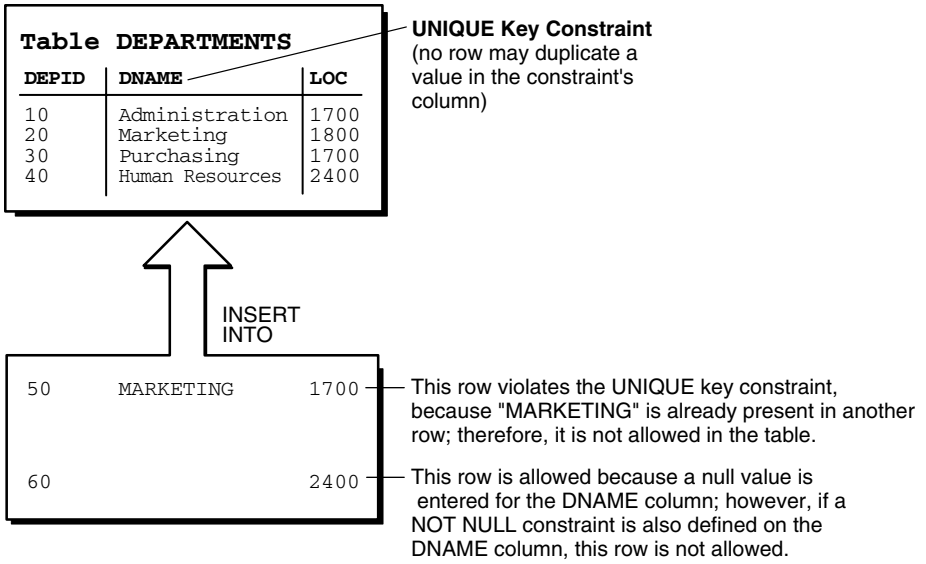

Note: You cannot have identical values in the non-null columns of a composite `UNIQUE` key constraint (`UNIQUE` key constraints allow `NULL` values).

Some examples of good unique keys include:

- An employee social security number (the primary key might be the employee number)
- A truck license plate number (the primary key might be the truck number)
- A customer phone number, consisting of the two columns `AREA_CODE` and `LOCAL_PHONE` (the primary key might be the customer number)
- A department name and location (the primary key might be the department number)

When to Use Constraints On Views

The constraints in this chapter apply to tables, not views.

Although you can declare constraints on views, such constraints do not help maintain data integrity. Instead, they are used to enable query rewrites on queries involving views, which helps performance with materialized views and other data warehousing features. Such constraints are always declared with the `DISABLE` keyword, and you cannot use the `VALIDATE` keyword. The constraints are never enforced, and there is no associated index.

See Also: *Oracle Database Data Warehousing Guide* for information about using constraints in data warehousing

Enforcing Referential Integrity with Constraints

Whenever two tables contain one or more common columns, Oracle Database can enforce the relationship between the two tables through a referential integrity constraint. Define a `PRIMARY` or `UNIQUE` key constraint on the column in the parent

table (the one that has the complete set of column values). Define a `FOREIGN KEY` constraint on the column in the child table (the one whose values must refer to existing values in the other table).

Note: In static data dictionary views `*_CONSTRAINTS`, a `FOREIGN KEY` constraint has `CONSTRAINT_TYPE` value `R` (for referential integrity).

See Also: "Defining Relationships Between Parent and Child Tables" on page 5-9 for information about defining additional constraints, including the foreign key

Figure 5–2 shows a foreign key defined on the department number. It guarantees that every value in this column must match a value in the primary key of the department table. This constraint prevents erroneous department numbers from getting into the employee table.

Foreign keys can be composed of multiple columns. Such a **composite foreign key** must reference a composite primary or unique key of the exact same structure, with the same number of columns and the same data types. Because composite primary and unique keys are limited to 32 columns, a composite foreign key is also limited to 32 columns.

Figure 5–2 Tables with FOREIGN KEY Constraints

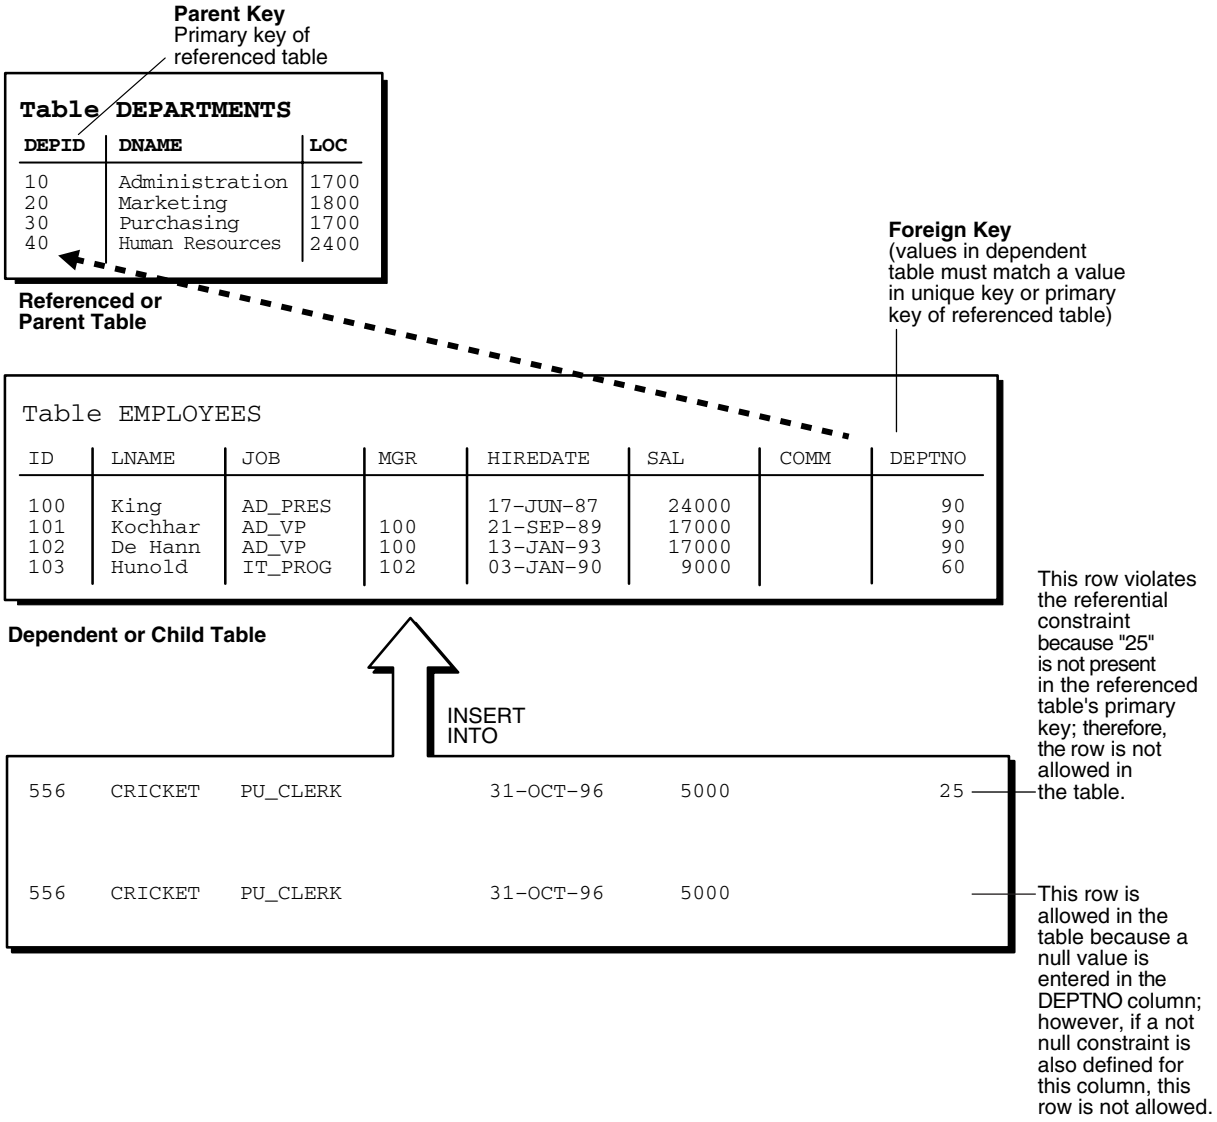

FOREIGN KEY Constraints and NULL Values

Foreign keys allow key values that are all `NULL`, even if there are no matching `PRIMARY` or `UNIQUE` keys.

- By default (without any `NOT NULL` or `CHECK` clauses), the `FOREIGN KEY` constraint enforces the **match none** rule for composite foreign keys in the ANSI/ISO standard.
- To enforce the **match full** rule for `NULL` values in composite foreign keys, which requires that all components of the key be `NULL` or all be non-null, define a `CHECK` constraint that allows only all nulls or all non-nulls in the composite foreign key. For example, with a composite key comprised of columns `A`, `B`, and `C`:

```
CHECK ((A IS NULL AND B IS NULL AND C IS NULL) OR
       (A IS NOT NULL AND B IS NOT NULL AND C IS NOT NULL))
```

- In general, it is not possible to use declarative referential integrity to enforce the **match partial** rule for `NULL` values in composite foreign keys, which requires the

non-null portions of the key to appear in the corresponding portions in the primary or unique key of a single row in the referenced table. You can often use triggers to handle this case, as described in *Oracle Database PL/SQL Language Reference*.

Defining Relationships Between Parent and Child Tables

Several relationships between parent and child tables can be determined by the other types of constraints defined on the foreign key in the child table.

No Constraints on the Foreign Key When no other constraints are defined on the foreign key, any number of rows in the child table can reference the same parent key value. This model allows nulls in the foreign key.

This model establishes a one-to-many relationship between the parent and foreign keys that allows undetermined values (nulls) in the foreign key. An example of such a relationship is shown in Figure 5–2 between the `employee` and `department` tables. Each department (parent key) has many employees (foreign key), and some employees might not be in a department (nulls in the foreign key).

NOT NULL Constraint on the Foreign Key When nulls are not allowed in a foreign key, each row in the child table must explicitly reference a value in the parent key because nulls are not allowed in the foreign key.

Any number of rows in the child table can reference the same parent key value, so this model establishes a one-to-many relationship between the parent and foreign keys. However, each row in the child table must have a reference to a parent key value; the absence of a value (a null) in the foreign key is not allowed. The same example in the previous section illustrates such a relationship. However, in this case, employees must have a reference to a specific department.

UNIQUE Constraint on the Foreign Key When a `UNIQUE` constraint is defined on the foreign key, only one row in the child table can reference a given parent key value. This model allows nulls in the foreign key.

This model establishes a one-to-one relationship between the parent and foreign keys that allows undetermined values (nulls) in the foreign key. For example, assume that the employee table had a column named `MEMBERNO`, referring to an employee membership number in the company insurance plan. Also, a table named `INSURANCE` has a primary key named `MEMBERNO`, and other columns of the table keep respective information relating to an employee insurance policy. The `MEMBERNO` in the employee table must be both a foreign key and a unique key:

- To enforce referential integrity rules between the `EMP_TAB` and `INSURANCE` tables (the `FOREIGN KEY` constraint)
- To guarantee that each employee has a unique membership number (the `UNIQUE` key constraint)

UNIQUE and NOT NULL Constraints on the Foreign Key When both `UNIQUE` and `NOT NULL` constraints are defined on the foreign key, only one row in the child table can reference a given parent key value, and because `NULL` values are not allowed in the foreign key, each row in the child table must explicitly reference a value in the parent key.

This model establishes a one-to-one relationship between the parent and foreign keys that does not allow undetermined values (nulls) in the foreign key. If you expand the previous example by adding a `NOT NULL` constraint on the `MEMBERNO` column of the

employee table, in addition to guaranteeing that each employee has a unique membership number, you also ensure that no undetermined values (nulls) are allowed in the MEMBERNO column of the employee table.

Rules for Multiple FOREIGN KEY Constraints

Oracle Database allows a column to be referenced by multiple FOREIGN KEY constraints; there is no limit on the number of dependent keys. This situation might be present if a single column is part of two different composite foreign keys.

Deferring Constraint Checks

When Oracle Database checks a constraint, it signals an error if the constraint is not satisfied. To defer checking constraints until the end of the current transaction, use the SET CONSTRAINTS statement.

> **Note:** You cannot use the SET CONSTRAINTS statement inside a trigger.

When deferring constraint checks:

- Select appropriate data.

  You might want to defer constraint checks on UNIQUE and FOREIGN keys if the data you are working with has any of these characteristics:

  - Tables are snapshots.
  - Some tables contain a large amount of data being manipulated by another application, which might not return the data in the same order.
- Update cascade operations on foreign keys.
- Ensure that constraints are deferrable.

  After identifying the appropriate tables, ensure that their FOREIGN, UNIQUE and PRIMARY key constraints are created DEFERRABLE.
- Within the application that manipulates the data, set all constraints deferred before you begin processing any data, as follows:

  ```
  SET CONSTRAINTS ALL DEFERRED;
  ```
- (Optional) Check for constraint violations immediately before committing the transaction.

  Immediately before the COMMIT statement, run the SET CONSTRAINTS ALL IMMEDIATE statement. If there are any problems with a constraint, this statement fails, and identifies the constraint that caused the error. If you commit while constraints are violated, the transaction rolls back and you get an error message.

In Example 5–2, the PRIMARY and FOREIGN keys of the table emp are created DEFERRABLE and then deferred.

Example 5–2 Deferring Constraint Checks

```
DROP TABLE dept;
CREATE TABLE dept (
  deptno NUMBER PRIMARY KEY,
  dname  VARCHAR2 (30)
);
```

```
DROP TABLE emp;
CREATE TABLE emp (
  empno  NUMBER,
  ename  VARCHAR2(30),
  deptno NUMBER,
  CONSTRAINT pk_emp_empno PRIMARY KEY (empno) DEFERRABLE,
  CONSTRAINT fk_emp_deptno FOREIGN KEY (deptno) REFERENCES dept(deptno) DEFERRABLE
);

INSERT INTO dept (deptno, dname) VALUES (10, 'Accounting');
INSERT INTO dept (deptno, dname) VALUES (20, 'SALES');

INSERT INTO emp (empno, ename, deptno) VALUES (1, 'Corleone', 10);
INSERT INTO emp (empno, ename, deptno) VALUES (2, 'Costanza', 20);
COMMIT;

SET CONSTRAINTS ALL DEFERRED;

UPDATE dept
SET deptno = deptno + 10
WHERE deptno = 20;
```

Query:

```
SELECT * from dept
ORDER BY deptno;
```

Result:

```
    DEPTNO DNAME
---------- ------------------------------
        10 Accounting
        30 SALES

2 rows selected.
```

Update:

```
UPDATE emp
SET deptno = deptno + 10
WHERE deptno = 20;
```

Result:

```
1 row updated.
```

Query:

```
SELECT * from emp
ORDER BY deptno;
```

Result:

```
     EMPNO ENAME                             DEPTNO
---------- ------------------------------ ----------
         1 Corleone                               10
         2 Costanza                               30

2 rows selected.
```

The `SET CONSTRAINTS` applies only to the current transaction, and its setting lasts for the duration of the transaction, or until another `SET CONSTRAINTS` statement resets the mode. The `ALTER SESSION SET CONSTRAINTS` statement applies only for the current session. The defaults specified when you create a constraint remain while the constraint exists.

> **See Also:** *Oracle Database SQL Language Reference* for more information about the `SET CONSTRAINTS` statement

Minimizing Space and Time Overhead for Indexes Associated with Constraints

When you create a `UNIQUE` or `PRIMARY` key, Oracle Database checks to see if an existing index enforces uniqueness for the constraint. If there is no such index, the database creates one.

When Oracle Database uses a unique index to enforce a constraint, and constraints associated with the unique index are dropped or disabled, the index is dropped. To preserve the statistics associated with the index (which would take a long time to re-create), specify the `KEEP INDEX` clause on the `DROP CONSTRAINT` statement.

While enabled foreign keys reference a `PRIMARY` or `UNIQUE` key, you cannot disable or drop the `PRIMARY` or `UNIQUE` key constraint or the index.

> **Note:** `UNIQUE` and `PRIMARY` keys with deferrable constraints must all use nonunique indexes.

To use existing indexes when creating unique and primary key constraints, include `USING INDEX` in the `CONSTRAINT` clause. For details and examples, see *Oracle Database SQL Language Reference*.

Guidelines for Indexing Foreign Keys

Index foreign keys unless the matching unique or primary key is never updated or deleted.

> **See Also:** *Oracle Database Concepts* for more information about indexing foreign keys

Referential Integrity in a Distributed Database

The declaration of a referential constraint cannot specify a foreign key that references a primary or unique key of a remote table.

However, you can maintain parent/child table relationships across nodes using triggers.

> **See Also:** *Oracle Database PL/SQL Language Reference* for more information about triggers that enforce referential integrity

Note: If you decide to define referential integrity across the nodes of a distributed database using triggers, be aware that network failures can make both the parent table and the child table inaccessible.

For example, assume that the child table is in the SALES database, and the parent table is in the HQ database.

If the network connection between the two databases fails, then some data manipulation language (DML) statements against the child table (those that insert rows or update a foreign key value) cannot proceed, because the referential integrity triggers must have access to the parent table in the HQ database.

When to Use CHECK Constraints

Use CHECK constraints when you must enforce integrity rules based on logical expressions, such as comparisons. Never use CHECK constraints when any of the other types of constraints can provide the necessary checking.

See Also: "Choosing Between CHECK and NOT NULL Constraints" on page 5-14

Examples of CHECK constraints include:

- A CHECK constraint on employee salaries so that no salary value is greater than 10000.
- A CHECK constraint on department locations so that only the locations "BOSTON", "NEW YORK", and "DALLAS" are allowed.
- A CHECK constraint on the salary and commissions columns to prevent the commission from being larger than the salary.

Restrictions on CHECK Constraints

A CHECK constraint requires that a condition be true or unknown for every row of the table. If a statement causes the condition to evaluate to false, then the statement is rolled back. The condition of a CHECK constraint has these limitations:

- The condition must be a boolean expression that can be evaluated using the values in the row being inserted or updated.
- The condition cannot contain subqueries or sequences.
- The condition cannot include the SYSDATE, UID, USER, or USERENV SQL functions.
- The condition cannot contain the pseudocolumns LEVEL or ROWNUM.
- The condition cannot contain the PRIOR operator.
- The condition cannot contain a user-defined function.

See Also:

- *Oracle Database SQL Language Reference* for information about the `LEVEL` pseudocolumn
- *Oracle Database SQL Language Reference* for information about the `ROWNUM` pseudocolumn
- *Oracle Database SQL Language Reference* for information about the `PRIOR` operator (used in hierarchical queries)

Designing CHECK Constraints

When using `CHECK` constraints, remember that a `CHECK` constraint is violated only if the condition evaluates to false; true and unknown values (such as comparisons with nulls) do not violate a check condition. Ensure that any `CHECK` constraint that you define is specific enough to enforce the rule.

For example, consider this `CHECK` constraint:

```
CHECK (Sal > 0 OR Comm >= 0)
```

At first glance, this rule may be interpreted as "do not allow a row in the employee table unless the employee salary is greater than zero or the employee commission is greater than or equal to zero." But if a row is inserted with a null salary, that row does not violate the `CHECK` constraint, regardless of whether the commission value is valid, because the entire check condition is evaluated as unknown. In this case, you can prevent such violations by placing `NOT NULL` constraints on both the `SAL` and `COMM` columns.

Note: If you are not sure when unknown values result in `NULL` conditions, review the truth tables for the logical conditions in *Oracle Database SQL Language Reference*

Rules for Multiple CHECK Constraints

A single column can have multiple `CHECK` constraints that reference the column in its definition. There is no limit to the number of `CHECK` constraints that can be defined that reference a column.

The order in which the constraints are evaluated is not defined, so be careful not to rely on the order or to define multiple constraints that conflict with each other.

Choosing Between CHECK and NOT NULL Constraints

According to the ANSI/ISO standard, a `NOT NULL` constraint is an example of a `CHECK` constraint, where the condition is:

```
CHECK (column_name IS NOT NULL)
```

Therefore, you can write `NOT NULL` constraints for a single column using either a `NOT NULL` constraint or a `CHECK` constraint. The `NOT NULL` constraint is easier to use than the `CHECK` constraint.

In the case where a composite key can allow only all nulls or all values, you must use a `CHECK` constraint. For example, this `CHECK` constraint allows a key value in the composite key made up of columns `C1` and `C2` to contain either all nulls or all values:

```
CHECK ((C1 IS NULL AND C2 IS NULL) OR (C1 IS NOT NULL AND C2 IS NOT NULL))
```

Examples of Defining Constraints

Example 5–3 and Example 5–4 show how to create simple constraints during the prototype phase of your database design. In these examples, each constraint is given a name. Naming the constraints prevents the database from creating multiple copies of the same constraint, with different system-generated names, if the data definition language (DDL) statement runs multiple times.

Example 5–3 creates tables and their constraints at the same time, using the `CREATE TABLE` statement.

Example 5–3 Defining Constraints with the CREATE TABLE Statement

```
DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3) CONSTRAINT pk_DeptTab_Deptno PRIMARY KEY,
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15),
  CONSTRAINT u_DeptTab_Dname_Loc UNIQUE (Dname, Loc),
  CONSTRAINT c_DeptTab_Loc
    CHECK (Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO')));

DROP TABLE EmpTab;
CREATE TABLE EmpTab (
  Empno    NUMBER(5) CONSTRAINT pk_EmpTab_Empno PRIMARY KEY,
  Ename    VARCHAR2(15) NOT NULL,
  Job      VARCHAR2(10),
  Mgr      NUMBER(5) CONSTRAINT r_EmpTab_Mgr REFERENCES EmpTab,
  Hiredate DATE,
  Sal      NUMBER(7,2),
  Comm     NUMBER(5,2),
  Deptno   NUMBER(3) NOT NULL
  CONSTRAINT r_EmpTab_DeptTab REFERENCES DeptTab ON DELETE CASCADE);
```

Example 5–4 creates constraints for existing tables, using the `ALTER TABLE` statement.

You cannot create a validated constraint on a table if the table contains rows that violate the constraint.

Example 5–4 Defining Constraints with the ALTER TABLE Statement

```
-- Create tables without constraints:

DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3),
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15)
);

DROP TABLE EmpTab;
CREATE TABLE EmpTab (
  Empno    NUMBER(5),
  Ename    VARCHAR2(15),
  Job      VARCHAR2(10),
  Mgr      NUMBER(5),
  Hiredate DATE,
  Sal      NUMBER(7,2),
  Comm     NUMBER(5,2),
  Deptno   NUMBER(3)
);
```

```
--Define constraints with the ALTER TABLE statement:

ALTER TABLE DeptTab
ADD CONSTRAINT pk_DeptTab_Deptno PRIMARY KEY (Deptno);

ALTER TABLE EmpTab
ADD CONSTRAINT fk_DeptTab_Deptno
FOREIGN KEY (Deptno) REFERENCES DeptTab;

ALTER TABLE EmpTab MODIFY (Ename VARCHAR2(15) NOT NULL);
```

> **See Also:** *Oracle Database Administrator's Guide* for information about creating and maintaining constraints for a large production database

Privileges Needed to Define Constraints

The creator of a constraint must have the ability to create tables (the `CREATE TABLE` or `CREATE ANY TABLE` system privilege), or the ability to alter the table (the `ALTER` object privilege for the table or the `ALTER ANY TABLE` system privilege) with the constraint. Additionally, `UNIQUE` and `PRIMARY KEY` constraints require that the owner of the table have either a quota for the tablespace that contains the associated index or the `UNLIMITED TABLESPACE` system privilege. `FOREIGN KEY` constraints also require some additional privileges.

> **See Also:** "Privileges Required to Create FOREIGN KEY Constraints" on page 5-23

Naming Constraints

Assign names to constraints `NOT NULL`, `UNIQUE`, `PRIMARY KEY`, `FOREIGN KEY`, and `CHECK` using the `CONSTRAINT` option of the constraint clause. This name must be unique among the constraints that you own. If you do not specify a constraint name, one is assigned automatically by Oracle Database.

Choosing your own name makes error messages for constraint violations more understandable, and prevents the creation of duplicate constraints with different names if the SQL statements are run more than once.

See the previous examples of the `CREATE TABLE` and `ALTER TABLE` statements for examples of the `CONSTRAINT` option of the `constraint` clause. The name of each constraint is included with other information about the constraint in the data dictionary.

> **See Also:** "Viewing Information About Constraints" on page 5-24 for examples of static data dictionary views

Enabling and Disabling Constraints

This section explains the mechanisms and procedures for manually enabling and disabling constraints.

enabled constraint. When a constraint is enabled, the corresponding rule is enforced on the data values in the associated columns. The definition of the constraint is stored in the data dictionary.

disabled constraint. When a constraint is disabled, the corresponding rule is not enforced. The definition of the constraint is still stored in the data dictionary.

An integrity constraint represents an assertion about the data in a database. This assertion is always true when the constraint is enabled. The assertion might not be true when the constraint is disabled, because data that violates the integrity constraint can be in the database.

Topics:

- Why Disable Constraints?
- Creating Enabled Constraints (Default)
- Creating Disabled Constraints
- Enabling Existing Constraints
- Disabling Existing Constraints
- Guidelines for Enabling and Disabling Key Constraints
- Fixing Constraint Exceptions

Why Disable Constraints?

During day-to-day operations, keep constraints enabled. In certain situations, temporarily disabling the constraints of a table makes sense for performance reasons. For example:

- When loading large amounts of data into a table using SQL*Loader
- When performing batch operations that make massive changes to a table (such as changing each employee number by adding 1000 to the existing number)
- When importing or exporting one table at a time

Temporarily turning off constraints can speed up these operations.

Creating Enabled Constraints (Default)

When you define an integrity constraint (using either `CREATE TABLE` or `ALTER TABLE`), Oracle Database enables the constraint by default. For code clarity, you can explicitly enable the constraint by including the `ENABLE` clause in its definition, as in Example 5–5.

Example 5–5 Creating Enabled Constraints

```
/* Use CREATE TABLE statement to create enabled constraint
   (ENABLE keyword is optional): */

DROP TABLE t1;
CREATE TABLE t1 (Empno NUMBER(5) PRIMARY KEY ENABLE);

/* Create table without constraint
   and then use ALTER TABLE statement to add enabled constraint
   (ENABLE keyword is optional): */

DROP TABLE t2;
CREATE TABLE t2 (Empno NUMBER(5));

ALTER TABLE t2 ADD PRIMARY KEY (Empno) ENABLE;
```

Include the `ENABLE` clause when defining a constraint for a table to be populated a row at a time by individual transactions. This ensures that data is always consistent, and reduces the performance overhead of each DML statement.

An ALTER TABLE statement that tries to enable an integrity constraint fails if an existing row of the table violates the integrity constraint. The statement rolls back and the constraint definition is neither stored nor enabled.

> **See Also:** "Fixing Constraint Exceptions" on page 5-19 for more information about rows that violate constraints

Creating Disabled Constraints

You define and disable an integrity constraint (using either CREATE TABLE or ALTER TABLE), by including the DISABLE clause in its definition, as in Example 5–6.

Example 5–6 Creating Disabled Constraints

```
/* Use CREATE TABLE statement to create disabled constraint */

DROP TABLE t1;
CREATE TABLE t1 (Empno NUMBER(5) PRIMARY KEY DISABLE);

/* Create table without constraint
   and then use ALTER TABLE statement to add disabled constraint */

DROP TABLE t2;
CREATE TABLE t2 (Empno NUMBER(5));

ALTER TABLE t2 ADD PRIMARY KEY (Empno) DISABLE;
```

Include the DISABLE clause when defining a constraint for a table to have large amounts of data inserted before anybody else accesses it, particularly if you must cleanse data after inserting it, or must fill empty columns with sequence numbers or parent/child relationships.

An ALTER TABLE statement that defines and disables a constraint never fails, because its rule is not enforced.

Enabling Existing Constraints

After you have cleansed the data and filled the empty columns, you can enable constraints that were disabled during data insertion.

To enable an existing constraint, use the ALTER TABLE statement with the ENABLE clause, as in Example 5–7.

Example 5–7 Enabling Existing Constraints

```
-- Create table with disabled constraints:

DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3) PRIMARY KEY DISABLE,
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15),
  CONSTRAINT uk_DeptTab_Dname_Loc UNIQUE (Dname, Loc) DISABLE,
  CONSTRAINT c_DeptTab_Loc
  CHECK (Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO')) DISABLE
);

-- Enable constraints:

 ALTER TABLE DeptTab
```

```
ENABLE PRIMARY KEY
ENABLE CONSTRAINT uk_DeptTab_Dname_Loc
ENABLE CONSTRAINT c_DeptTab_Loc;
```

An ALTER TABLE statement that attempts to enable an integrity constraint fails if any of the table rows violate the integrity constraint. The statement is rolled back and the constraint is not enabled.

> **See Also:** "Fixing Constraint Exceptions" on page 5-19 for more information about rows that violate constraints

Disabling Existing Constraints

If you must perform a large insert or update when a table contains data, you can temporarily disable constraints to improve performance of the bulk operation.

To disable an existing constraint, use the ALTER TABLE statement with the DISABLE clause, as in Example 5–8.

Example 5–8 Disabling Existing Constraints

```
-- Create table with enabled constraints:

DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3) PRIMARY KEY ENABLE,
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15),
  CONSTRAINT uk_DeptTab_Dname_Loc UNIQUE (Dname, Loc) ENABLE,
  CONSTRAINT c_DeptTab_Loc
  CHECK (Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO')) ENABLE
);

-- Disable constraints:

ALTER TABLE DeptTab
DISABLE PRIMARY KEY
DISABLE CONSTRAINT uk_DeptTab_Dname_Loc
DISABLE CONSTRAINT c_DeptTab_Loc;
```

Guidelines for Enabling and Disabling Key Constraints

When enabling or disabling UNIQUE, PRIMARY KEY, and FOREIGN KEY constraints, be aware of several important issues and prerequisites. UNIQUE key and PRIMARY KEY constraints are usually managed by the database administrator.

> **See Also:** *Oracle Database Administrator's Guide* and "Managing FOREIGN KEY Constraints" on page 5-22

Fixing Constraint Exceptions

If a row of a table disobeys an integrity constraint, then this row is in violation of the constraint and is called an **exception** to the constraint. If any exceptions exist, then the constraint cannot be enabled. The rows that violate the constraint must be updated or deleted before the constraint can be enabled.

You can identify exceptions for a specific integrity constraint as you try to enable the constraint.

> **See Also:** "Fixing Constraint Exceptions" on page 5-19 for more information about this procedure

When you try to create or enable a constraint, and the statement fails because integrity constraint exceptions exist, the statement is rolled back. You cannot enable the constraint until all exceptions are either updated or deleted. To determine which rows violate the integrity constraint, include the `EXCEPTIONS` option in the `ENABLE` clause of a `CREATE TABLE` or `ALTER TABLE` statement.

> **See Also:** *Oracle Database Administrator's Guide* for more information about responding to constraint exceptions

Modifying Constraints

Starting with Oracle8*i*, you can modify an existing constraint with the `MODIFY CONSTRAINT` clause, as in Example 5–9.

> **See Also:** *Oracle Database SQL Language Reference* for information about the parameters you can modify

Example 5–9 Modifying Constraints

```
/* Create & then modify a CHECK constraint: */

DROP TABLE X1Tab;
CREATE TABLE X1Tab (
  a1 NUMBER
  CONSTRAINT c_X1Tab_a1 CHECK (a1>3)
  DEFERRABLE DISABLE
);

ALTER TABLE X1Tab
MODIFY CONSTRAINT c_X1Tab_a1 ENABLE;

ALTER TABLE X1Tab
MODIFY CONSTRAINT c_X1Tab_a1 RELY;

ALTER TABLE X1Tab
MODIFY CONSTRAINT c_X1Tab_a1 INITIALLY DEFERRED;

ALTER TABLE X1Tab
MODIFY CONSTRAINT c_X1Tab_a1 ENABLE NOVALIDATE;

/* Create & then modify a PRIMARY KEY constraint: */

DROP TABLE t1;
CREATE TABLE t1 (a1 INT, b1 INT);

ALTER TABLE t1
ADD CONSTRAINT pk_t1_a1 PRIMARY KEY(a1) DISABLE;

ALTER TABLE t1
MODIFY PRIMARY KEY INITIALLY IMMEDIATE
USING INDEX PCTFREE = 30 ENABLE NOVALIDATE;

ALTER TABLE t1
MODIFY PRIMARY KEY ENABLE NOVALIDATE;
```

Renaming Constraints

One property of a constraint that you can modify is its name. Situations in which you would rename a constraint include:

- You want to clone a table and its constraints.

  Constraint names must be unique, even across multiple schemas. Therefore, the constraints in the original table cannot have the same names as those in the cloned table.

- You created a constraint with a default system-generated name, and now you want to give it a name that is easy to remember, so that you can easily enable and disable it.

Example 5–10 shows how to find the system-generated name of a constraint and change it.

Example 5–10 Renaming a Constraint

```
DROP TABLE T;
CREATE TABLE T (
  C1 NUMBER PRIMARY KEY,
  C2 NUMBER
);
```

Query:

```
SELECT CONSTRAINT_NAME FROM USER_CONSTRAINTS
WHERE TABLE_NAME = 'T'
AND CONSTRAINT_TYPE = 'P';
```

Result (system-generated name of constraint name varies):

```
CONSTRAINT_NAME
------------------------------
SYS_C0013059

1 row selected.
```

Rename constraint from name reported in preceding query to `T_C1_PK`:

```
ALTER TABLE T
RENAME CONSTRAINT SYS_C0013059
TO T_C1_PK;
```

Query:

```
SELECT CONSTRAINT_NAME FROM USER_CONSTRAINTS
WHERE TABLE_NAME = 'T'
AND CONSTRAINT_TYPE = 'P';
```

Result:

```
CONSTRAINT_NAME
------------------------------
T_C1_PK

1 row selected.
```

Dropping Constraints

You can drop a constraint using the DROP clause of the ALTER TABLE statement. Situations in which you would drop a constraint include:

- The constraint enforces a rule that is no longer true.
- The constraint is no longer needed.

To drop a constraint and all other integrity constraints that depend on it, specify CASCADE.

Example 5–11 Dropping Constraints

```
-- Create table with constraints:

DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3) PRIMARY KEY,
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15),
  CONSTRAINT uk_DeptTab_Dname_Loc UNIQUE (Dname, Loc),
  CONSTRAINT c_DeptTab_Loc
  CHECK (Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO'))
);

-- Drop constraints:

ALTER TABLE DeptTab
DROP PRIMARY KEY
DROP CONSTRAINT uk_DeptTab_Dname_Loc
DROP CONSTRAINT c_DeptTab_Loc;
```

When dropping UNIQUE, PRIMARY KEY, and FOREIGN KEY constraints, be aware of several important issues and prerequisites. UNIQUE and PRIMARY KEY constraints are usually managed by the database administrator.

See Also:

- *Oracle Database SQL Language Reference* for more information about the DROP clause of the ALTER TABLE statement.
- *Oracle Database Administrator's Guide* for more information about dropping constraints.
- *Oracle Database SQL Language Reference* for information about the CASCADE CONSTRAINTS clause of the DROP TABLE statement, which drops all referential integrity constraints that refer to primary and unique keys in the dropped table

Managing FOREIGN KEY Constraints

FOREIGN KEY constraints enforce relationships between columns in different tables. Therefore, they cannot be enabled if the constraint of the referenced primary or unique key is not present or not enabled.

Data Types and Names for Foreign Key Columns

You must use the same data type for corresponding columns in the dependent and referenced tables. The column names need not match.

Limit on Columns in Composite Foreign Keys

Because foreign keys reference primary and unique keys of the parent table, and PRIMARY KEY and UNIQUE key constraints are enforced using indexes, composite foreign keys are limited to 32 columns.

Foreign Key References Primary Key by Default

If the column list is not included in the REFERENCES option when defining a FOREIGN KEY constraint (single column or composite), then Oracle Database assumes that you intend to reference the primary key of the specified table. Alternatively, you can explicitly specify the column(s) to reference in the parent table within parentheses. Oracle Database automatically checks to verify that this column list references a primary or unique key of the parent table. If it does not, then an informative error is returned.

Privileges Required to Create FOREIGN KEY Constraints

To create a FOREIGN KEY constraint, the creator of the constraint must have privileged access to the parent and child tables.

- **Parent Table** The creator of the referential integrity constraint must own the parent table or have REFERENCES object privileges on the columns that constitute the parent key of the parent table.
- **Child Table** The creator of the referential integrity constraint must have the ability to create tables (that is, the CREATE TABLE or CREATE ANY TABLE system privilege) or the ability to alter the child table (that is, the ALTER object privilege for the child table or the ALTER ANY TABLE system privilege).

In both cases, necessary privileges cannot be obtained through a role; they must be explicitly granted to the creator of the constraint.

These restrictions allow:

- The owner of the child table to explicitly decide which constraints are enforced and which other users can create constraints
- The owner of the parent table to explicitly decide if foreign keys can depend on the primary and unique keys in her tables

Choosing How Foreign Keys Enforce Referential Integrity

Oracle Database allows different types of referential integrity actions to be enforced, as specified with the definition of a FOREIGN KEY constraint:

- **Prevent Delete or Update of Parent Key** The default setting prevents the deletion or update of a parent key if there is a row in the child table that references the key. For example:

  ```
  CREATE TABLE Emp_tab (
  FOREIGN KEY (Deptno) REFERENCES Dept_tab);
  ```

- **Delete Child Rows When Parent Key Deleted** The ON DELETE CASCADE action allows parent key data that is referenced from the child table to be deleted, but not updated. When data in the parent key is deleted, all rows in the child table that depend on the deleted parent key values are also deleted. To specify this referential action, include the ON DELETE CASCADE option in the definition of the FOREIGN KEY constraint. For example:

  ```
  CREATE TABLE Emp_tab (
  ```

```
FOREIGN KEY (Deptno) REFERENCES Dept_tab
ON DELETE CASCADE);
```

- **Set Foreign Keys to Null When Parent Key Deleted** The `ON DELETE SET NULL` action allows data that references the parent key to be deleted, but not updated. When referenced data in the parent key is deleted, all rows in the child table that depend on those parent key values have their foreign keys set to null. To specify this referential action, include the `ON DELETE SET NULL` option in the definition of the `FOREIGN KEY` constraint. For example:

```
CREATE TABLE Emp_tab (
FOREIGN KEY (Deptno) REFERENCES Dept_tab
ON DELETE SET NULL);
```

Viewing Information About Constraints

To find the names of constraints, what columns they affect, and other information to help you manage them, query the static data dictionary views `*_CONSTRAINTS` and `*_CONS_COLUMNS`, as in Example 5–12.

> **See Also:** *Oracle Database Reference* for information about `*_CONSTRAINTS` and `*_CONS_COLUMNS`

Example 5–12 Viewing Information About Constraints

```
DROP TABLE DeptTab;
CREATE TABLE DeptTab (
  Deptno  NUMBER(3) PRIMARY KEY,
  Dname   VARCHAR2(15),
  Loc     VARCHAR2(15),
  CONSTRAINT uk_DeptTab_Dname_Loc UNIQUE (Dname, Loc),
  CONSTRAINT c_DeptTab_Loc
  CHECK (Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO'))
);

DROP TABLE EmpTab;
CREATE TABLE EmpTab (
  Empno    NUMBER(5) PRIMARY KEY,
  Ename    VARCHAR2(15) NOT NULL,
  Job      VARCHAR2(10),
  Mgr      NUMBER(5) CONSTRAINT r_EmpTab_Mgr
             REFERENCES EmpTab ON DELETE CASCADE,
  Hiredate DATE,
  Sal      NUMBER(7,2),
  Comm     NUMBER(5,2),
  Deptno   NUMBER(3) NOT NULL
  CONSTRAINT r_EmpTab_Deptno REFERENCES DeptTab
);

-- Format columns (optional):

COLUMN CONSTRAINT_NAME   FORMAT A20;
COLUMN CONSTRAINT_TYPE   FORMAT A4 HEADING 'TYPE';
COLUMN TABLE_NAME        FORMAT A10;
COLUMN R_CONSTRAINT_NAME FORMAT A17;
COLUMN SEARCH_CONDITION  FORMAT A40;
COLUMN COLUMN_NAME       FORMAT A12;
```

List accessible constraints in DeptTab and EmpTab:

```
SELECT CONSTRAINT_NAME, CONSTRAINT_TYPE, TABLE_NAME, R_CONSTRAINT_NAME
FROM USER_CONSTRAINTS
WHERE (TABLE_NAME = 'DEPTTAB' OR TABLE_NAME = 'EMPTAB')
ORDER BY CONSTRAINT_NAME;
```

Result:

```
CONSTRAINT_NAME      TYPE TABLE_NAME R_CONSTRAINT_NAME
-------------------- ---- ---------- -----------------
C_DEPTTAB_LOC        C    DEPTTAB
R_EMPTAB_DEPTNO      R    EMPTAB     SYS_C006286
R_EMPTAB_MGR         R    EMPTAB     SYS_C006290
SYS_C006286          P    DEPTTAB
SYS_C006288          C    EMPTAB
SYS_C006289          C    EMPTAB
SYS_C006290          P    EMPTAB
UK_DEPTTAB_DNAME_LOC U    DEPTTAB

8 rows selected.
```

Distinguish between `NOT NULL` and `CHECK` constraints in `DeptTab` and `EmpTab`:

```
SELECT CONSTRAINT_NAME, SEARCH_CONDITION
FROM USER_CONSTRAINTS
WHERE (TABLE_NAME = 'DEPTTAB' OR TABLE_NAME = 'EMPTAB')
AND CONSTRAINT_TYPE = 'C'
ORDER BY CONSTRAINT_NAME;
```

Result:

```
CONSTRAINT_NAME      SEARCH_CONDITION
-------------------- ----------------------------------------
C_DEPTTAB_LOC        Loc IN ('NEW YORK', 'BOSTON', 'CHICAGO')
SYS_C006288          "ENAME" IS NOT NULL
SYS_C006289          "DEPTNO" IS NOT NULL

3 rows selected.
```

For `DeptTab` and `EmpTab`, list columns that constitute constraints:

```
SELECT CONSTRAINT_NAME, TABLE_NAME, COLUMN_NAME
FROM USER_CONS_COLUMNS
WHERE (TABLE_NAME = 'DEPTTAB' OR TABLE_NAME = 'EMPTAB')
ORDER BY CONSTRAINT_NAME;
```

Result:

```
CONSTRAINT_NAME      TABLE_NAME COLUMN_NAME
-------------------- ---------- ------------
C_DEPTTAB_LOC        DEPTTAB    LOC
R_EMPTAB_DEPTNO      EMPTAB     DEPTNO
R_EMPTAB_MGR         EMPTAB     MGR
SYS_C006286          DEPTTAB    DEPTNO
SYS_C006288          EMPTAB     ENAME
SYS_C006289          EMPTAB     DEPTNO
SYS_C006290          EMPTAB     EMPNO
UK_DEPTTAB_DNAME_LOC DEPTTAB    LOC
UK_DEPTTAB_DNAME_LOC DEPTTAB    DNAME

9 rows selected.
```

Note that:

- Some constraint names are user specified (such as UK_DEPTTAB_DNAME_LOC), while others are system specified (such as SYS_C006290).
- Each constraint type is denoted with a different character in the CONSTRAINT_TYPE column. This table summarizes the characters used for each constraint type:

Constraint Type	Character
PRIMARY KEY	P
UNIQUE KEY	U
FOREIGN KEY	R
CHECK, NOT NULL	C

Note: An additional constraint type is indicated by the character "V" in the CONSTRAINT_TYPE column. This constraint type corresponds to constraints created using the WITH CHECK OPTION for views.

These constraints are explicitly listed in the SEARCH_CONDITION column:

- NOT NULL constraints
- The conditions for user-defined CHECK constraints

Part II

PL/SQL for Application Developers

This part presents information that application developers need about PL/SQL, the Oracle procedural extension of SQL.

Chapters:

- Chapter 6, "Coding PL/SQL Subprograms and Packages"
- Chapter 7, "Using PL/Scope"
- Chapter 8, "Using the PL/SQL Hierarchical Profiler"
- Chapter 9, "Developing PL/SQL Web Applications"
- Chapter 10, "Developing PL/SQL Server Pages (PSP)"
- Chapter 11, "Using Continuous Query Notification (CQN)"

> **See Also:** *Oracle Database PL/SQL Language Reference* for a complete description of PL/SQL

6

Coding PL/SQL Subprograms and Packages

This chapter describes some procedural capabilities of Oracle Database for application development.

Topics:

- Overview of PL/SQL Units
- Compiling PL/SQL Subprograms for Native Execution
- Cursor Variables
- Handling PL/SQL Compile-Time Errors
- Handling Run-Time PL/SQL Errors
- Debugging Stored Subprograms
- Invoking Stored Subprograms
- Invoking Remote Subprograms
- Invoking Stored PL/SQL Functions from SQL Statements
- Returning Large Amounts of Data from a Function
- Coding Your Own Aggregate Functions

See Also:

- *Oracle Database PL/SQL Language Reference* for more information about PL/SQL subprograms
- *Oracle Database PL/SQL Language Reference* for more information about PL/SQL packages
- *Oracle Database Performance Tuning Guide* for information about application tracing tools, which can help you find problems in PL/SQL code

Overview of PL/SQL Units

PL/SQL is a modern, block-structured programming language. It provides several features that make developing powerful database applications very convenient. For example, PL/SQL provides procedural constructs, such as loops and conditional statements, that are not available in standard SQL.

You can directly enter SQL data manipulation language (DML) statements inside PL/SQL blocks, and you can use subprograms supplied by Oracle to perform data definition language (DDL) statements.

PL/SQL code runs on the server, so using PL/SQL lets you centralize significant parts of your database applications for increased maintainability and security. It also enables you to achieve a significant reduction of network overhead in client/server applications.

> **Note:** Some Oracle tools, such as Oracle Forms, contain a PL/SQL engine that lets you run PL/SQL locally.

You can even use PL/SQL for some database applications instead of 3GL programs that use embedded SQL or Oracle Call Interface (OCI).

PL/SQL units include:

- Anonymous Blocks
- Stored PL/SQL Units
- Triggers

See Also:

- *Oracle Database PL/SQL Language Reference* for syntax and examples of operations on PL/SQL packages
- *Oracle Database PL/SQL Packages and Types Reference* for information about the PL/SQL packages that come with Oracle Database
- "Dependencies Among Local and Remote Database Procedures" on page 18-11 for information about dependencies among stored PL/SQL units

Anonymous Blocks

An anonymous block is a PL/SQL unit that has no name. An anonymous block consists of an optional declarative part, an executable part, and one or more optional exception handlers.

The declarative part declares PL/SQL variables, exceptions, and cursors. The executable part contains PL/SQL code and SQL statements, and can contain nested blocks.

Exception handlers contain code that is invoked when the exception is raised, either as a predefined PL/SQL exception (such as `NO_DATA_FOUND` or `ZERO_DIVIDE`) or as an exception that you define.

Anonymous blocks are usually used interactively from a tool, such as SQL*Plus, or in a precompiler, OCI, or SQL*Module application. They are usually used to invoke stored subprograms or to open cursor variables.

The anonymous block in Example 6–1 uses the `DBMS_OUTPUT` package to print the names of all employees in the `HR.EMPLOYEES` table who are in department 20.

Example 6–1 Anonymous Block

```
DECLARE
  last_name  VARCHAR2(10);
  cursor     c1 IS
               SELECT LAST_NAME
               FROM EMPLOYEES
               WHERE DEPARTMENT_ID = 20
```

```
              ORDER BY LAST_NAME;
BEGIN
  OPEN c1;
  LOOP
    FETCH c1 INTO last_name;
    EXIT WHEN c1%NOTFOUND;
    DBMS_OUTPUT.PUT_LINE(last_name);
  END LOOP;
END;
/
```

Result:

```
Fay
Hartstein
```

Exceptions let you handle Oracle Database error conditions with PL/SQL program logic, enabling your application to prevent the server from issuing an error that can cause the client application to end. The anonymous block in Example 6–2 handles the predefined Oracle Database exception `NO_DATA_FOUND` (which results in ORA-01403 if not handled).

Example 6–2 Anonymous Block with Exception Handler for Predefined Error

```
DECLARE
  Emp_number  INTEGER := 9999
  Emp_name    VARCHAR2(10);
BEGIN
  SELECT LAST_NAME INTO Emp_name
    FROM EMPLOYEES
      WHERE EMPLOYEE_ID = Emp_number;
  DBMS_OUTPUT.PUT_LINE('Employee name is ' || Emp_name);
EXCEPTION
  WHEN NO_DATA_FOUND THEN
    DBMS_OUTPUT.PUT_LINE('No such employee: ' || Emp_number);
END;
/
```

Result:

```
No such employee: 9999
```

You can also define your own exceptions; that is, you can declare them in the declaration part of a block and define them in the exception part of the block, as in Example 6–3.

Example 6–3 Anonymous Block with Exception Handler for User-Defined Exception

```
DECLARE
  Emp_name            VARCHAR2(10);
  Emp_number          INTEGER;
  Empno_out_of_range  EXCEPTION;
BEGIN
  Emp_number := 10001;
  IF Emp_number > 9999 OR Emp_number < 1000 THEN
    RAISE Empno_out_of_range;
  ELSE
    SELECT LAST_NAME INTO Emp_name
    FROM EMPLOYEES
    WHERE EMPLOYEE_ID = Emp_number;
    DBMS_OUTPUT.PUT_LINE('Employee name is ' || Emp_name);
 END IF;
EXCEPTION
```

```
  WHEN Empno_out_of_range THEN
    DBMS_OUTPUT.PUT_LINE('Employee number ' || Emp_number ||
      ' is out of range.');
END;
/
```

Result:

```
Employee number 10001 is out of range.
```

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for complete information about the `DBMS_OUTPUT` package
- *Oracle Database PL/SQL Language Reference* and "Handling Run-Time PL/SQL Errors" on page 6-23
- "Cursor Variables" on page 6-19

Stored PL/SQL Units

A stored PL/SQL unit is a subprogram (procedure or function) or package that:

- Has a name.
- Can take parameters, and can return values.
- Is stored in the data dictionary.
- Can be invoked by many users.

If a subprogram belongs to a package, it is called a **package subprogram**; if not, it is called a **standalone subprogram**.

Topics:

- Naming Subprograms
- Subprogram Parameters
- Creating Subprograms
- Altering Subprograms
- Dropping Subprograms and Packages
- External Subprograms
- PL/SQL Function Result Cache
- PL/SQL Packages
- PL/SQL Object Size Limits
- Creating Packages
- Naming Packages and Package Objects
- Package Invalidations and Session State
- Packages Supplied with Oracle Database
- Overview of Bulk Binding
- When to Use Bulk Binds

Naming Subprograms

Because a subprogram is stored in the database, it must be named. This distinguishes it from other stored subprograms and makes it possible for applications to invoke it. Each publicly-visible subprogram in a schema must have a unique name, and the name must be a legal PL/SQL identifier.

> **Note:** If you plan to invoke a stored subprogram using a stub generated by SQL*Module, then the stored subprogram name must also be a legal identifier in the invoking host 3GL language, such as Ada or C.

Subprogram Parameters

Stored subprograms can take parameters. In the procedure in Example 6–4, the department number is an input parameter that is used when the parameterized cursor `c1` is opened.

Example 6–4 Stored Procedure with Parameters

```
CREATE OR REPLACE PROCEDURE get_emp_names (
  dept_num IN NUMBER
)
IS
  emp_name  VARCHAR2(10);
  CURSOR    c1 (dept_num NUMBER) IS
                SELECT LAST_NAME FROM EMPLOYEES
                WHERE DEPARTMENT_ID = dept_num;
BEGIN
  OPEN c1(dept_num);
  LOOP
    FETCH c1 INTO emp_name;
    EXIT WHEN C1%NOTFOUND;
    DBMS_OUTPUT.PUT_LINE(emp_name);
  END LOOP;
  CLOSE c1;
END;
/
```

The formal parameters of a subprogram have three major attributes, described in Table 6–1.

Table 6–1 Attributes of Subprogram Parameters

Parameter Attribute	Description
Name	This must be a legal PL/SQL identifier.
Mode	This indicates whether the parameter is an input-only parameter (`IN`), an output-only parameter (`OUT`), or is both an input and an output parameter (`IN OUT`). If the mode is not specified, then `IN` is assumed.
Data Type	This is a standard PL/SQL data type.

Topics:

- Parameter Modes
- Parameter Data Types
- %TYPE and %ROWTYPE Attributes

- Passing Composite Variables as Parameters
- Initial Parameter Values

Parameter Modes Parameter modes define the action of formal parameters. You can use the three parameter modes, `IN` (the default), `OUT`, and `IN OUT`, with any subprogram. Avoid using the `OUT` and `IN OUT` modes with functions. Good programming practice dictates that a function returns a single value and does not change the values of variables that are not local to the subprogram.

Table 6–2 summarizes the information about parameter modes.

Table 6–2 Parameter Modes

IN	OUT	IN OUT
The default.	Must be specified.	Must be specified.
Passes values to a subprogram.	Returns values to the caller.	Passes initial values to a subprogram; returns updated values to the caller.
Formal parameter acts like a constant.	Formal parameter acts like an uninitialized variable.	Formal parameter acts like an initialized variable.
Formal parameter cannot be assigned a value.	Formal parameter cannot be used in an expression; must be assigned a value.	Formal parameter must be assigned a value.
Actual parameter can be a constant, initialized variable, literal, or expression.	Actual parameter must be a variable.	Actual parameter must be a variable.

> **See Also:** *Oracle Database PL/SQL Language Reference* for details about parameter modes

Parameter Data Types The data type of a formal parameter consists of one of these:

- An unconstrained type name, such as `NUMBER` or `VARCHAR2`.
- A type that is constrained using the `%TYPE` or `%ROWTYPE` attributes.

> **Note:** Numerically constrained types such as `NUMBER(2)` or `VARCHAR2(20)` are not allowed in a parameter list.

%TYPE and %ROWTYPE Attributes Use the type attributes `%TYPE` and `%ROWTYPE` to constrain the parameter. For example, the procedure heading in Example 6–4 can be written as follows:

```
PROCEDURE get_emp_names(dept_num IN EMPLOYEES.DEPARTMENT_ID%TYPE)
```

This gives the `dept_num` parameter the same data type as the `DEPARTMENT_ID` column in the `EMPLOYEES` table. The column and table must be available when a declaration using `%TYPE` (or `%ROWTYPE`) is elaborated.

Using `%TYPE` is recommended, because if the type of the column in the table changes, it is not necessary to change the application code.

If the `get_emp_names` procedure is part of a package, you can use previously-declared public (package) variables to constrain its parameter data types. For example:

```
dept_number  NUMBER(2);
...
PROCEDURE get_emp_names(dept_num IN dept_number%TYPE);
```

Use the %ROWTYPE attribute to create a record that contains all the columns of the specified table. The procedure in Example 6–5 returns all the columns of the EMPLOYEES table in a PL/SQL record for the given employee ID.

Example 6–5 %TYPE and %ROWTYPE Attributes

```
CREATE OR REPLACE PROCEDURE get_emp_rec (
  emp_number  IN  EMPLOYEES.EMPLOYEE_ID%TYPE,
  emp_info    OUT EMPLOYEES%ROWTYPE
)
IS
BEGIN
  SELECT * INTO emp_info
  FROM EMPLOYEES
  WHERE EMPLOYEE_ID = emp_number;
END;
/
```

Invoke procedure from PL/SQL block:

```
DECLARE
  emp_row  EMPLOYEES%ROWTYPE;
BEGIN
  get_emp_rec(206, emp_row);
  DBMS_OUTPUT.PUT('EMPLOYEE_ID: ' || emp_row.EMPLOYEE_ID);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('FIRST_NAME: ' || emp_row.FIRST_NAME);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('LAST_NAME: ' || emp_row.LAST_NAME);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('EMAIL: ' || emp_row.EMAIL);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('PHONE_NUMBER: ' || emp_row.PHONE_NUMBER);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('HIRE_DATE: ' || emp_row.HIRE_DATE);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('JOB_ID: ' || emp_row.JOB_ID);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('SALARY: ' || emp_row.SALARY);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('COMMISSION_PCT: ' || emp_row.COMMISSION_PCT);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('MANAGER_ID: ' || emp_row.MANAGER_ID);
  DBMS_OUTPUT.NEW_LINE;
  DBMS_OUTPUT.PUT('DEPARTMENT_ID: ' || emp_row.DEPARTMENT_ID);
  DBMS_OUTPUT.NEW_LINE;
END;
/
```

Result:

```
EMPLOYEE_ID: 206
FIRST_NAME: William
LAST_NAME: Gietz
EMAIL: WGIETZ
PHONE_NUMBER: 515.123.8181
HIRE_DATE: 07-JUN-02
JOB_ID: AC_ACCOUNT
```

```
SALARY: 8300
COMMISSION_PCT:
MANAGER_ID: 205
DEPARTMENT_ID: 110
```

Stored functions can return values that are declared using `%ROWTYPE`. For example:

```
FUNCTION get_emp_rec (dept_num IN EMPLOYEES.DEPARTMENT_ID%TYPE)
   RETURN EMPLOYEES%ROWTYPE IS ...
```

Passing Composite Variables as Parameters You can pass PL/SQL composite variables (collections and records) as parameters to stored subprograms.

If the subprogram is remote, you must create a redundant loop-back DBLINK, so that when the remote subprogram compiles, the type checker that verifies the source uses the same definition of the user-defined composite variable type as the invoker uses.

Initial Parameter Values Parameters can take initial values. Use either the assignment operator or the `DEFAULT` keyword to give a parameter an initial value. For example, these are equivalent:

```
PROCEDURE Get_emp_names (Dept_num IN NUMBER := 20) IS ...
PROCEDURE Get_emp_names (Dept_num IN NUMBER DEFAULT) IS ...
```

When a parameter takes an initial value, it can be omitted from the actual parameter list when you invoke the subprogram. When you do specify the parameter value on the invocation, it overrides the initial value.

> **Note:** Unlike in an anonymous PL/SQL block, you do not use the keyword `DECLARE` before the declarations of variables, cursors, and exceptions in a stored subprogram. In fact, it is an error to use it.

Creating Subprograms

Use a text editor to write the subprogram. Then, using an interactive tool such as SQL*Plus, load the text file containing the procedure by entering:

```
@get_emp
```

This loads the procedure into the current schema from the `get_emp.sql` file (`.sql` is the default file extension). The slash (/) after the code is not part of the code, it only activates the loading of the procedure.

> **Caution:** When developing a subprogram, it is usually preferable to use the statement `CREATE OR REPLACE PROCEDURE` or `CREATE OR REPLACE FUNCTION`. This statement replaces any previous version of that subprogram in the same schema with the newer version, but without warning.

You can use either the keyword `IS` or `AS` after the subprogram parameter list.

See Also:

- *Oracle Database SQL Language Reference* for the syntax of the `CREATE FUNCTION` statement
- *Oracle Database SQL Language Reference* for the syntax of the `CREATE PROCEDURE` statement

Privileges Needed

To create a subprogram, a package specification, or a package body, you must meet these prerequisites:

- You must have the `CREATE PROCEDURE` system privilege to create a subprogram or package in your schema, or the `CREATE ANY PROCEDURE` system privilege to create a subprogram or package in another user's schema. In either case, the package body must be created in the same schema as the package.

  Note: To create without errors (to compile the subprogram or package successfully) requires these additional privileges:

  - The owner of the subprogram or package must be explicitly granted the necessary object privileges for all objects referenced within the body of the code.
  - The owner cannot obtain required privileges through roles.

If the privileges of the owner of a subprogram or package change, then the subprogram must be reauthenticated before it is run. If a necessary privilege to a referenced object is revoked from the owner of the subprogram or package, then the subprogram cannot be run.

The `EXECUTE` privilege on a subprogram gives a user the right to run a subprogram owned by another user. Privileged users run the subprogram under the security domain of the owner of the subprogram. Therefore, users need not be granted the privileges to the objects referenced by a subprogram. This allows for more disciplined and efficient security strategies with database applications and their users. Furthermore, all subprograms and packages are stored in the data dictionary (in the `SYSTEM` tablespace). No quota controls the amount of space available to a user who creates subprograms and packages.

Note: Package creation requires a sort. The user creating the package must be able to create a sort segment in the temporary tablespace with which the user is associated.

Altering Subprograms

To alter a subprogram, you must first drop it using the `DROP PROCEDURE` or `DROP FUNCTION` statement, then re-create it using the `CREATE PROCEDURE` or `CREATE FUNCTION` statement. Alternatively, use the `CREATE OR REPLACE PROCEDURE` or `CREATE OR REPLACE FUNCTION` statement, which first drops the subprogram if it exists, then re-creates it as specified.

Caution: The subprogram is dropped without warning.

Dropping Subprograms and Packages

A standalone subprogram, a standalone function, a package body, or an entire package can be dropped using the SQL statements `DROP PROCEDURE`, `DROP FUNCTION`, `DROP PACKAGE BODY`, and `DROP PACKAGE`, respectively. A `DROP PACKAGE` statement drops both the specification and body of a package.

This statement drops the `Old_sal_raise` procedure in your schema:

```
DROP PROCEDURE Old_sal_raise;
```

Privileges Needed

To drop a subprogram or package, the subprogram or package must be in your schema, or you must have the `DROP ANY PROCEDURE` privilege. An individual subprogram within a package cannot be dropped; the containing package specification and body must be re-created without the subprograms to be dropped.

External Subprograms

A PL/SQL subprogram running on an Oracle Database instance can invoke an external subprogram written in a third-generation language (3GL). The 3GL subprogram runs in a separate address space from that of the database.

> **See Also:** Chapter 14, "Developing Applications with Multiple Programming Languages," for information about external subprograms

PL/SQL Function Result Cache

Using the PL/SQL function result cache can save significant space and time. Each time a **result-cached** PL/SQL function is invoked with different parameter values, those parameters and their result are stored in the cache. Subsequently, when the same function is invoked with the same parameter values, the result is retrieved from the cache, instead of being recomputed. Because the cache is stored in a shared global area (SGA), it is available to any session that runs your application.

If a database object that was used to compute a cached result is updated, the cached result becomes invalid and must be recomputed.

The best candidates for result-caching are functions that are invoked frequently but depend on information that changes infrequently or never.

For more information about the PL/SQL function result cache, see *Oracle Database PL/SQL Language Reference*.

PL/SQL Packages

A **package** is a collection of related program objects (for example, subprogram, variables, constants, cursors, and exceptions) stored as a unit in the database.

Using packages is an alternative to creating subprograms as standalone schema objects. Packages have many advantages over standalone subprograms. For example, they:

- Let you organize your application development more efficiently.
- Let you grant privileges more efficiently.
- Let you modify package objects without recompiling dependent schema objects.
- Enable Oracle Database to read multiple package objects into memory at once.

- Can contain global variables and cursors that are available to all subprograms in the package.
- Let you **overload** subprograms. Overloading a subprogram means creating multiple subprograms with the same name in the same package, each taking arguments of different number or data type.

  See Also: *Oracle Database PL/SQL Language Reference* for more information about subprogram name overloading

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

Example 6–6 creates a package that contains one stored function and two stored procedures, and then invokes a procedure.

Example 6–6 Creating PL/SQL Package and Invoking Package Subprogram

```
-- Sequence that package function needs:

CREATE SEQUENCE emp_sequence
START WITH 8000
INCREMENT BY 10;

-- Package specification:

CREATE or REPLACE PACKAGE employee_management IS
  FUNCTION hire_emp (
  firstname  VARCHAR2,
  lastname   VARCHAR2,
  email      VARCHAR2,
  phone      VARCHAR2,
  hiredate   DATE,
  job        VARCHAR2,
  sal        NUMBER,
  comm       NUMBER,
  mgr        NUMBER,
  deptno     NUMBER
) RETURN NUMBER;

 PROCEDURE fire_emp(
    emp_id IN NUMBER
 );

 PROCEDURE sal_raise (
    emp_id IN NUMBER,
    sal_incr IN NUMBER
 );
END employee_management;
/

-- Package body:

CREATE or REPLACE PACKAGE BODY employee_management IS
  FUNCTION hire_emp (
    firstname  VARCHAR2,
    lastname   VARCHAR2,
    email      VARCHAR2,
```

```
   phone      VARCHAR2,
   hiredate   DATE,
   job        VARCHAR2,
   sal        NUMBER,
   comm       NUMBER,
   mgr        NUMBER,
   deptno     NUMBER
 ) RETURN NUMBER
IS
  new_empno  NUMBER(10);
BEGIN
  new_empno := emp_sequence.NEXTVAL;

   INSERT INTO EMPLOYEES (
     employee_id,
     first_name,
     last_name,
     email,
     phone_number,
     hire_date,
     job_id,
     salary,
     commission_pct,
     manager_id,
     department_id
   )
   VALUES (
     new_empno,
     firstname,
     lastname,
     email,
     phone,
     hiredate,
     job,
     sal,
     comm,
     mgr,
     deptno
   );

  RETURN (new_empno);
END hire_emp;

PROCEDURE fire_emp (
  emp_id IN NUMBER
) IS
BEGIN
  DELETE FROM EMPLOYEES
  WHERE EMPLOYEE_ID = emp_id;

  IF SQL%NOTFOUND THEN
    raise_application_error(
      -20011,
      'Invalid Employee Number: ' || TO_CHAR(Emp_id)
    );
  END IF;
END fire_emp;

PROCEDURE sal_raise (
   emp_id IN NUMBER,
```

```
    sal_incr IN NUMBER
  ) IS
  BEGIN
    UPDATE EMPLOYEES
    SET SALARY = SALARY + sal_incr
    WHERE EMPLOYEE_ID = emp_id;

    IF SQL%NOTFOUND THEN
      raise_application_error(
        -20011,
        'Invalid Employee Number: ' || TO_CHAR(Emp_id)
      );
    END IF;
  END sal_raise;
END employee_management;
/
```

Invoke package procedures:

```
DECLARE
  empno  NUMBER(6);
  sal    NUMBER(6);
  temp   NUMBER(6);
BEGIN
  empno := employee_management.hire_emp(
           'John',
           'Doe',
           'john.doe@company.com',
           '555-0100',
           '20-SEP-07',
           'ST_CLERK',
           2500,
           0,
           100,
           20);

  DBMS_OUTPUT.PUT_LINE('New employee ID is ' || TO_CHAR(empno));
END;
/
```

PL/SQL Object Size Limits

The size limit for PL/SQL stored database objects such as subprograms, triggers, and packages is the size of the **Descriptive Intermediate Attributed Notation for Ada (DIANA)** code in the shared pool in bytes. The Linux and UNIX limit on the size of the flattened DIANA/code size is 64K but the limit might be 32K on desktop platforms.

The most closely related number that a user can access is the `PARSED_SIZE` in the static data dictionary view `*_OBJECT_SIZE`. That gives the size of the DIANA in bytes as stored in the `SYS.IDL_xxx$` tables. This is not the size in the shared pool. The size of the DIANA part of PL/SQL code (used during compilation) is significantly larger in the shared pool than it is in the system table.

Creating Packages

Each part of a package is created with a different statement. Create the package specification using the `CREATE PACKAGE` statement. The `CREATE PACKAGE` statement declares public package objects.

To create a package body, use the `CREATE PACKAGE BODY` statement. The `CREATE PACKAGE BODY` statement defines the procedural code of the public subprograms declared in the package specification.

You can also define private, or local, package subprograms, and variables in a package body. These objects can only be accessed by other subprograms in the body of the same package. They are not visible to external users, regardless of the privileges they hold.

It is often more convenient to add the `OR REPLACE` clause in the `CREATE PACKAGE` or `CREATE PACKAGE BODY` statements when you are first developing your application. The effect of this option is to drop the package or the package body without warning. The `CREATE` statements are:

```
CREATE OR REPLACE PACKAGE Package_name AS ...
```

and

```
CREATE OR REPLACE PACKAGE BODY Package_name AS ...
```

Creating Package Objects The body of a package can contain:

- Subprograms declared in the package specification.
- Definitions of cursors declared in the package specification.
- Local subprograms, not declared in the package specification.
- Local variables.

Subprograms, cursors, and variables that are declared in the package specification are **global**. They can be invoked, or used, by external users that have `EXECUTE` permission for the package or that have `EXECUTE ANY PROCEDURE` privileges.

When you create the package body, ensure that each subprogram that you define in the body has the same parameters, by name, data type, and mode, as the declaration in the package specification. For functions in the package body, the parameters and the return type must agree in name and type.

Privileges to Needed to Create or Drop Packages The privileges required to create or drop a package specification or package body are the same as those required to create or drop a standalone subprogram. See "Creating Subprograms" on page 6-8 and "Dropping Subprograms and Packages" on page 6-10.

Naming Packages and Package Objects

The names of a package and all public objects in the package must be unique within a given schema. The package specification and its body must have the same name. All package constructs must have unique names within the scope of the package, unless overloading of subprogram names is desired.

Package Invalidations and Session State

Each session that references a package object has its own instance of the corresponding package, including persistent state for any public and private variables, cursors, and constants. If any of the session's instantiated packages (specification or body) are invalidated, then all package instances in the session are invalidated and recompiled. Therefore, the session state is lost for all package instances in the session.

When a package in a given session is invalidated, the session receives ORA-04068 the first time it attempts to use any object of the invalid package instance. The second time a session makes such a package call, the package is reinstantiated for the session

without error. However, if you handle this error in your application, be aware of the following:

- For optimal performance, Oracle Database returns this error message only when the package state is discarded. When a subprogram in one package invokes a subprogram in another package, the session state is lost for both packages.
- If a server session traps ORA-04068, then ORA-04068 is not raised for the client session. Therefore, when the client session attempts to use an object in the package, the package is not reinstantiated. To reinstantiate the package, the client session must either reconnect to the database or recompile the package.

In Example 6–7, the `RAISE` statement raises the current exception, ORA-04068, which is the cause of the exception being handled, ORA-06508. ORA-04068 is not trapped.

Example 6–7 Raising ORA-04068

```
PROCEDURE p IS
  package_exception EXCEPTION;
  PRAGMA EXCEPTION_INIT (package_exception, -6508);
BEGIN
 ...
EXCEPTION
  WHEN package_exception THEN
    RAISE;
END;
/
```

In Example 6–8, the `RAISE` statement raises the exception ORA-20001 in response to ORA-06508, instead of the current exception, ORA-04068. ORA-04068 is trapped. When this happens, the ORA-04068 error is masked, which stops the package from being reinstantiated.

Example 6–8 Trapping ORA-04068

```
PROCEDURE p IS
  package_exception EXCEPTION;
  other_exception   EXCEPTION;
  PRAGMA EXCEPTION_INIT (package_exception, -6508);
  PRAGMA EXCEPTION_INIT (other_exception, -20001);
BEGIN
 ...
EXCEPTION
  WHEN package_exception THEN
    ...
    RAISE other_exception;
END;
/
```

In most production environments, DDL operations that can cause invalidations are usually performed during inactive working hours; therefore, this situation might not be a problem for end-user applications. However, if package invalidations are common in your system during working hours, then you might want to code your applications to handle this error when package calls are made.

Packages Supplied with Oracle Database

There are many packages provided with Oracle Database, either to extend the functionality of the database or to give PL/SQL access to SQL features. You can invoke these packages from your application.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for an overview of these Oracle Database packages

Overview of Bulk Binding

Oracle Database uses two engines to run PL/SQL blocks and subprograms. The PL/SQL engine runs procedural statements, while the SQL engine runs SQL statements. During execution, every SQL statement causes a context switch between the two engines, resulting in performance overhead.

Performance can be improved substantially by minimizing the number of context switches required to run a particular block or subprogram. When a SQL statement runs inside a loop that uses collection elements as bind variables, the large number of context switches required by the block can cause poor performance. Collections include:

- Varrays
- Nested tables
- Index-by tables
- Host arrays

Binding is the assignment of values to PL/SQL variables in SQL statements. **Bulk binding** is binding an entire collection at once. Bulk binds pass the entire collection back and forth between the two engines in a single operation.

Typically, using bulk binds improves performance for SQL statements that affect four or more database rows. The more rows affected by a SQL statement, the greater the performance gain from bulk binds.

> **Note:** This section provides an overview of bulk binds to help you decide whether to use them in your PL/SQL applications. For detailed information about using bulk binds, including ways to handle exceptions that occur in the middle of a bulk bind operation, see *Oracle Database PL/SQL Language Reference.*
>
> Parallel DML statements are disabled with bulk binds.

When to Use Bulk Binds

Consider using bulk binds to improve the performance of:

- DML Statements that Reference Collections
- SELECT Statements that Reference Collections
- FOR Loops that Reference Collections and Return DML

DML Statements that Reference Collections A bulk bind, which uses the `FORALL` keyword, can improve the performance of `INSERT`, `UPDATE`, or `DELETE` statements that reference collection elements.

The PL/SQL block in Example 6–9 increases the salary for employees whose manager's ID number is 7902, 7698, or 7839, with and without bulk binds. Without bulk bind, PL/SQL sends a SQL statement to the SQL engine for each updated employee, leading to context switches that slow performance.

Example 6–9 DML Statements that Reference Collections

```
DECLARE
```

```
  TYPE numlist IS VARRAY (100) OF NUMBER;
  id NUMLIST := NUMLIST(7902, 7698, 7839);
BEGIN
  -- Efficient method, using bulk bind:

  FORALL i IN id.FIRST..id.LAST
  UPDATE EMPLOYEES
  SET SALARY = 1.1 * SALARY
  WHERE MANAGER_ID = id(i);

 -- Slower method:

 FOR i IN id.FIRST..id.LAST LOOP
    UPDATE EMPLOYEES
    SET SALARY = 1.1 * SALARY
    WHERE MANAGER_ID = id(i);
 END LOOP;
END;
/
```

SELECT Statements that Reference Collections The `BULK COLLECT INTO` clause can improve the performance of queries that reference collections. You can use `BULK COLLECT INTO` with tables of scalar values, or tables of `%TYPE` values.

The PL/SQL block in Example 6–10 queries multiple values into PL/SQL tables, with and without bulk binds. Without bulk bind, PL/SQL sends a SQL statement to the SQL engine for each selected employee, leading to context switches that slow performance.

Example 6–10 SELECT Statements that Reference Collections

```
DECLARE
  TYPE var_tab IS TABLE OF VARCHAR2(20)
  INDEX BY PLS_INTEGER;

  empno    VAR_TAB;
  ename    VAR_TAB;
  counter  NUMBER;

  CURSOR c IS
    SELECT EMPLOYEE_ID, LAST_NAME
    FROM EMPLOYEES
    WHERE MANAGER_ID = 7698;
BEGIN
 -- Efficient method, using bulk bind:

 SELECT EMPLOYEE_ID, LAST_NAME BULK COLLECT
 INTO empno, ename
 FROM EMPLOYEES
 WHERE MANAGER_ID = 7698;

 -- Slower method:

 counter := 1;

 FOR rec IN c LOOP
    empno(counter) := rec.EMPLOYEE_ID;
    ename(counter) := rec.LAST_NAME;
    counter := counter + 1;
 END LOOP;
```

```
END;
/
```

FOR Loops that Reference Collections and Return DML You can use the FORALL keyword with the BULK COLLECT INTO keywords to improve the performance of FOR loops that reference collections and return DML.

The PL/SQL block in Example 6–11 updates the EMPLOYEES table by computing bonuses for a collection of employees. Then it returns the bonuses in a column called bonus_list_inst. The actions are performed with and without bulk binds. Without bulk bind, PL/SQL sends a SQL statement to the SQL engine for each updated employee, leading to context switches that slow performance.

Example 6–11 FOR Loops that Reference Collections and Return DML

```
DECLARE
  TYPE emp_list IS VARRAY(100) OF EMPLOYEES.EMPLOYEE_ID%TYPE;
  empids emp_list := emp_list(182, 187, 193, 200, 204, 206);

  TYPE bonus_list IS TABLE OF EMPLOYEES.SALARY%TYPE;
  bonus_list_inst  bonus_list;

BEGIN
  -- Efficient method, using bulk bind:

 FORALL i IN empids.FIRST..empids.LAST
 UPDATE EMPLOYEES
 SET SALARY = 0.1 * SALARY
 WHERE EMPLOYEE_ID = empids(i)
 RETURNING SALARY BULK COLLECT INTO bonus_list_inst;

 -- Slower method:

 FOR i IN empids.FIRST..empids.LAST LOOP
   UPDATE EMPLOYEES
   SET SALARY = 0.1 * SALARY
   WHERE EMPLOYEE_ID = empids(i)
   RETURNING SALARY INTO bonus_list_inst(i);
 END LOOP;
END;
/
```

Triggers

A trigger is a special kind of PL/SQL anonymous block. You can define triggers to fire before or after SQL statements, either on a statement level or for each row that is affected. You can also define INSTEAD OF triggers or system triggers (triggers on DATABASE and SCHEMA).

> **See Also:** *Oracle Database PL/SQL Language Reference* for more information about triggers

Compiling PL/SQL Subprograms for Native Execution

You can speed up PL/SQL subprograms by compiling them into native code residing in shared libraries.

You can use native compilation with both the supplied packages and the subprograms you write yourself. Subprograms compiled this way work in all server environments,

such as the shared server configuration (formerly known as multithreaded server) and Oracle Real Application Clusters (Oracle RAC).

This technique is most effective for computation-intensive subprograms that do not spend much time running SQL, because it can do little to speed up SQL statements invoked from these subprograms.

With Java, you can use the `ncomp` tool to compile your own packages and classes.

See Also:

- *Oracle Database PL/SQL Language Reference* for details on PL/SQL native compilation
- *Oracle Database Java Developer's Guide* for details on Java native compilation

Cursor Variables

A cursor is a static object; a cursor variable is a pointer to a cursor. Because cursor variables are pointers, they can be passed and returned as parameters to subprograms. A cursor variable can also refer to different cursors in its lifetime.

Additional advantages of cursor variables include:

- Encapsulation

  Queries are centralized in the stored subprogram that opens the cursor variable.

- Easy maintenance

  If you must change the cursor, then you only make the change in the stored subprogram, not in each application.

- Convenient security

  The user of the application is the user name used when the application connects to the server. The user must have `EXECUTE` permission on the stored subprogram that opens the cursor. But, the user need not have `READ` permission on the tables used in the query. Use this capability to limit access to the columns in the table and access to other stored subprograms.

  See Also: *Oracle Database PL/SQL Language Reference* for more information about cursor variables

Topics:

- Declaring and Opening Cursor Variables
- Examples of Cursor Variables

Declaring and Opening Cursor Variables

Memory is usually allocated for a cursor variable in the client application using the appropriate `ALLOCATE` statement. In Pro*C, use the `EXEC SQL ALLOCATE` *`cursor_name`* statement. In OCI, use the Cursor Data Area.

You can also use cursor variables in applications that run entirely in a single server session. You can declare cursor variables in PL/SQL subprograms, open them, and use them as parameters for other PL/SQL subprograms.

Examples of Cursor Variables

This section has these examples of cursor variable usage in PL/SQL:

- Example 6–12, "Fetching Data with Cursor Variable"
- Example 6–13, "Cursor Variable with Discriminator"

> **See Also:** For additional cursor variable examples that use programmatic interfaces:
>
> - *Pro*COBOL Programmer's Guide*
> - *Oracle Call Interface Programmer's Guide*

Example 6–12 creates a package that defines a PL/SQL cursor variable type and two procedures, and then invokes the procedures from a PL/SQL block. The first procedure opens a cursor variable using a bind variable in the `WHERE` clause. The second procedure uses a cursor variable to fetch rows from the `EMPLOYEES` table.

Example 6–12 Fetching Data with Cursor Variable

```
CREATE OR REPLACE PACKAGE emp_data AS
  TYPE emp_val_cv_type IS REF CURSOR
  RETURN EMPLOYEES%ROWTYPE;

  PROCEDURE open_emp_cv (
    emp_cv       IN OUT emp_val_cv_type,
    dept_number  IN     EMPLOYEES.DEPARTMENT_ID%TYPE
  );

 PROCEDURE fetch_emp_data (
   emp_cv   IN  emp_val_cv_type,
   emp_row  OUT EMPLOYEES%ROWTYPE
 );
END emp_data;
/
CREATE OR REPLACE PACKAGE BODY emp_data AS
  PROCEDURE open_emp_cv (
    emp_cv       IN OUT emp_val_cv_type,
    dept_number  IN     EMPLOYEES.DEPARTMENT_ID%TYPE
  )
  IS
  BEGIN
    OPEN emp_cv FOR
    SELECT * FROM EMPLOYEES
    WHERE DEPARTMENT_ID = dept_number
    ORDER BY last_name;
  END open_emp_cv;

  PROCEDURE fetch_emp_data (
    emp_cv   IN  emp_val_cv_type,
    emp_row  OUT EMPLOYEES%ROWTYPE
  )
  IS
  BEGIN
    FETCH emp_cv INTO emp_row;
  END fetch_emp_data;
END emp_data;
/
```

Invoke package procedures:

```
DECLARE
  emp_curs     emp_data.emp_val_cv_type;
  dept_number  EMPLOYEES.DEPARTMENT_ID%TYPE;
  emp_row      EMPLOYEES%ROWTYPE;

BEGIN
  dept_number := 20;

  -- Open cursor, using variable:

 emp_data.open_emp_cv(emp_curs, dept_number);

 -- Fetch and display data:

 LOOP
   emp_data.fetch_emp_data(emp_curs, emp_row);
   EXIT WHEN emp_curs%NOTFOUND;
   DBMS_OUTPUT.PUT(emp_row.LAST_NAME || '  ');
   DBMS_OUTPUT.PUT_LINE(emp_row.SALARY);
 END LOOP;
END;
/
```

In Example 6–13, the procedure opens a cursor variable for either the EMPLOYEES table or the DEPARTMENTS table, depending on the value of the parameter discrim. The anonymous block invokes the procedure to open the cursor variable for the EMPLOYEES table, but fetches from the DEPARTMENTS table, which raises the predefined exception ROWTYPE_MISMATCH.

Example 6–13 Cursor Variable with Discriminator

```
CREATE OR REPLACE PACKAGE emp_dept_data AS
  TYPE cv_type IS REF CURSOR;

  PROCEDURE open_cv (
    cv       IN OUT cv_type,
    discrim  IN     POSITIVE
  );
  END emp_dept_data;
/

CREATE OR REPLACE PACKAGE BODY emp_dept_data AS
  PROCEDURE open_cv (
    cv      IN OUT cv_type,
    discrim IN     POSITIVE) IS
  BEGIN
    IF discrim = 1 THEN
    OPEN cv FOR
      SELECT * FROM EMPLOYEES ORDER BY employee_id;
    ELSIF discrim = 2 THEN
      OPEN cv FOR
        SELECT * FROM DEPARTMENTS ORDER BY department_id;
    END IF;
  END open_cv;
END emp_dept_data;
/
```

Invoke procedure open_cv from anonymous block:

```
DECLARE
  emp_rec   EMPLOYEES%ROWTYPE;
  dept_rec  DEPARTMENTS%ROWTYPE;
  cv        Emp_dept_data.CV_TYPE;
BEGIN
  emp_dept_data.open_cv(cv, 1);  -- Open cv for EMPLOYEES fetch.
  FETCH cv INTO dept_rec;        -- Fetch from DEPARTMENTS.
  DBMS_OUTPUT.PUT(dept_rec.DEPARTMENT_ID);
  DBMS_OUTPUT.PUT_LINE('  ' || dept_rec.LOCATION_ID);
EXCEPTION
  WHEN ROWTYPE_MISMATCH THEN
     BEGIN
       DBMS_OUTPUT.PUT_LINE
         ('Row type mismatch, fetching EMPLOYEES data ...');
       FETCH cv INTO emp_rec;
       DBMS_OUTPUT.PUT(emp_rec.DEPARTMENT_ID);
       DBMS_OUTPUT.PUT_LINE('  ' || emp_rec.LAST_NAME);
     END;
END;
/
```

Result:

```
Row type mismatch, fetching EMPLOYEES data ...
90  King
```

Handling PL/SQL Compile-Time Errors

To list compile-time errors, query the static data dictionary view `*_ERRORS`. From these views, you can retrieve original source code. The error text associated with the compilation of a subprogram is updated when the subprogram is replaced, and it is deleted when the subprogram is dropped.

SQL*Plus issues a warning message for compile-time errors, but for more information about them, you must use the command `SHOW ERRORS`.

> **Note:** Before issuing the `SHOW ERRORS` statement, use the `SET LINESIZE` statement to get long lines on output. The value 132 is usually a good choice. For example:
>
> ```
> SET LINESIZE 132
> ```

Example 6–14 has two compile-time errors: `WHER` should be `WHERE`, and `END` should be followed by a semicolon. `SHOW ERRORS` shows the line, column, and description of each error.

Example 6–14 Compile-Time Errors

```
CREATE OR REPLACE PROCEDURE fire_emp (
  emp_id NUMBER
) AS
BEGIN
  DELETE FROM EMPLOYEES
  WHER EMPLOYEE_ID = Emp_id;
END
/
```

Result:

```
Warning: Procedure created with compilation errors.
```

Command:

```
SHOW ERRORS;
```

Result:

```
Errors for PROCEDURE FIRE_EMP:

LINE/COL ERROR
-------- -----------------------------------------------------------------
5/3      PL/SQL: SQL Statement ignored
6/8    PL/SQL: ORA-00933: SQL command not properly ended
7/3    PLS-00103: Encountered the symbol "end-of-file" when expecting
         one of the following:
         ; <an identifier> <a double-quoted delimited-identifier>
         current delete exists prior <a single-quoted SQL string>
         The symbol ";" was substituted for "end-of-file" to continue.
```

See Also:

- *Oracle Database Reference* for more information about the static data dictionary view `*_SOURCE`
- *SQL*Plus User's Guide and Reference* for more information about the `SHOW ERRORS` statement

Handling Run-Time PL/SQL Errors

Oracle Database allows user-defined errors in PL/SQL code to be handled so that user-specified error numbers and messages are returned to the client application, which can handle the error.

User-specified error messages are returned using the `RAISE_APPLICATION_ERROR` procedure. For example:

```
RAISE_APPLICATION_ERROR(error_number, 'text', keep_error_stack)
```

This procedure stops subprogram execution, rolls back any effects of the subprogram, and returns a user-specified error number and message (unless the error is trapped by an exception handler). `error_number` must be in the range of -20000 to -20999.

Use error number -20000 as a generic number for messages where it is important to relay information to the user, but having a unique error number is not required. `Text` must be a character expression, 2 KB or less (longer messages are ignored). To add the error to errors on the stack, set `Keep_error_stack` to `TRUE`; to replace the existing errors, set it to `FALSE` (the default).

Note: Some Oracle Database packages, such as `DBMS_OUTPUT`, `DBMS_DESCRIBE`, and `DBMS_ALERT`, use application error numbers in the range -20000 to -20005. See the descriptions of these packages for more information.

The `RAISE_APPLICATION_ERROR` procedure is often used in exception handlers or in the logic of PL/SQL code. For example, this exception handler selects the string for the associated user-defined error message and invokes the `RAISE_APPLICATION_ERROR` procedure:

```
...
```

```
WHEN NO_DATA_FOUND THEN
   SELECT Error_string INTO Message
   FROM Error_table,
   V$NLS_PARAMETERS V
   WHERE Error_number = -20101 AND Lang = v.value AND
      v.parameter = "NLS_LANGUAGE";
   Raise_application_error(-20101, Message);
...
```

Topics:

- Declaring Exceptions and Exception Handlers
- Unhandled Exceptions
- Handling Errors in Distributed Queries
- Handling Errors in Remote Subprograms

Declaring Exceptions and Exception Handlers

User-defined exceptions are explicitly defined and raised within the PL/SQL block, to process errors specific to the application. When an exception is raised, the usual execution of the PL/SQL block stops, and an exception handler is invoked. Specific exception handlers can be written to handle any internal or user-defined exception.

Application code can check for a condition that requires special attention using an `IF` statement. If there is an error condition, then two options are available:

- Enter a `RAISE` statement that names the appropriate exception. A `RAISE` statement stops the execution of the subprogram, and control passes to an exception handler (if any).
- Invoke the `RAISE_APPLICATION_ERROR` procedure to return a user-specified error number and message.

You can also define an exception handler to handle user-specified error messages. For example, Figure 6–1 shows:

- An exception and associated exception handler in a subprogram
- A conditional statement that checks for an error (such as transferring funds not available) and enters a user-specified error number and message within a trigger
- How user-specified error numbers are returned to the invoking environment (in this case, a subprogram), and how that application can define an exception that corresponds to the user-specified error number

Declare a user-defined exception in a subprogram or package body (private exceptions), or in the specification of a package (public exceptions). Define an exception handler in the body of a subprogram (standalone or package).

Figure 6–1 Exceptions and User-Defined Errors

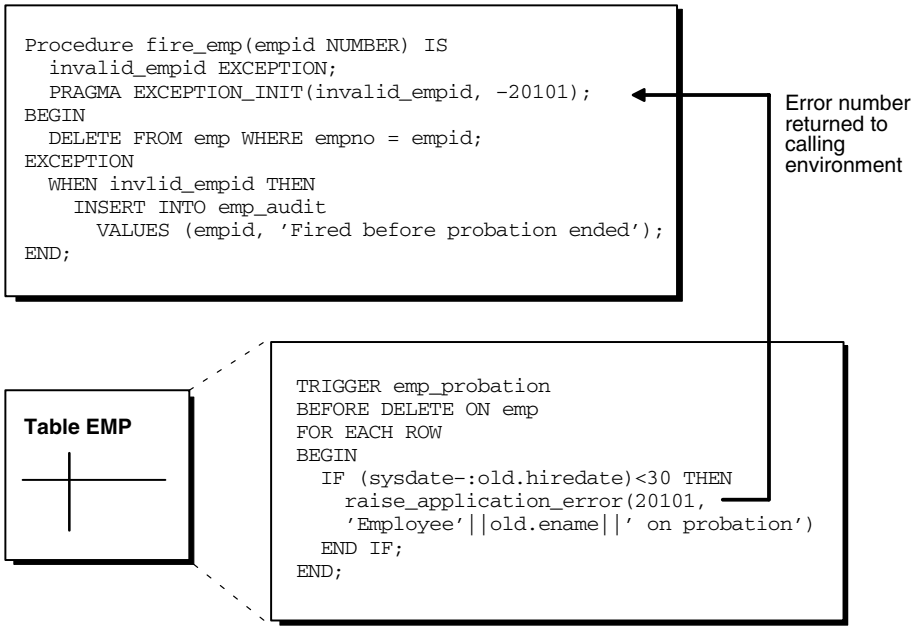

Unhandled Exceptions

In database PL/SQL units, an unhandled user-error condition or internal error condition that is not trapped by an appropriate exception handler causes the implicit rollback of the program unit. If the program unit includes a `COMMIT` statement before the point at which the unhandled exception is observed, then the implicit rollback of the program unit can only be completed back to the previous `COMMIT`.

Additionally, unhandled exceptions in database-stored PL/SQL units propagate back to client-side applications that invoke the containing program unit. In such an application, only the application program unit invocation is rolled back (not the entire application program unit), because it is submitted to the database as a SQL statement.

If unhandled exceptions in database PL/SQL units are propagated back to database applications, modify the database PL/SQL code to handle the exceptions. Your application can also trap for unhandled exceptions when invoking database program units and handle such errors appropriately.

Handling Errors in Distributed Queries

You can use a trigger or a stored subprogram to create a distributed query. This distributed query is decomposed by the local Oracle Database instance into a corresponding number of remote queries, which are sent to the remote nodes for execution. The remote nodes run the queries and send the results back to the local node. The local node then performs any necessary post-processing and returns the results to the user or application.

If a portion of a distributed statement fails, possibly from a constraint violation, then Oracle Database returns ORA-02055. Subsequent statements, or subprogram invocations, return ORA-02067 until a rollback or a rollback to savepoint is entered.

Design your application to check for any returned error messages that indicates that a portion of the distributed update has failed. If you detect a failure, rollback the entire transaction (or rollback to a savepoint) before allowing the application to proceed.

Handling Errors in Remote Subprograms

When a subprogram is run locally or at a remote location, these types of exceptions can occur:

- PL/SQL user-defined exceptions, which must be declared using the keyword `EXCEPTION`
- PL/SQL predefined exceptions, such as `NO_DATA_FOUND`
- SQL errors, such as ORA-00900
- Application exceptions, which are generated using the `RAISE_APPLICATION_ERROR` procedure.

When using local subprograms, all of these messages can be trapped by writing an exception handler, such as:

```
EXCEPTION
    WHEN ZERO_DIVIDE THEN
    /* Handle the exception */
```

The `WHEN` clause requires an exception name. If the exception that is raised does not have a name, such as those generated with `RAISE_APPLICATION_ERROR`, then one can be assigned using `PRAGMA_EXCEPTION_INIT`. For example:

```
DECLARE
    ...
    Null_salary EXCEPTION;
    PRAGMA EXCEPTION_INIT(Null_salary, -20101);
BEGIN
    ...
    RAISE_APPLICATION_ERROR(-20101, 'salary is missing');
    ...
EXCEPTION
    WHEN Null_salary THEN
        ...
```

When invoking a remote subprogram, exceptions are also handled by creating a local exception handler. The remote subprogram must return an error number to the local invoking subprogram, which then handles the exception, as shown in the previous example. Because PL/SQL user-defined exceptions always return ORA-06510 to the local subprogram, these exceptions cannot be handled. All other remote exceptions can be handled in the same manner as local exceptions.

Debugging Stored Subprograms

Compiling a stored subprogram involves fixing any syntax errors in the code. You might need to do additional debugging to ensure that the subprogram works correctly, performs well, and recovers from errors. Such debugging might involve:

- Adding extra output statements to verify execution progress and check data values at certain points within the subprogram.
- Running a separate debugger to analyze execution in greater detail.

Topics:

- PL/Scope
- PL/SQL Hierarchical Profiler
- Oracle JDeveloper

- DBMS_OUTPUT Package
- Privileges for Debugging PL/SQL and Java Stored Subprograms
- Writing Low-Level Debugging Code
- DBMS_DEBUG_JDWP Package
- DBMS_DEBUG Package

PL/Scope

PL/Scope is a compiler-driven tool that collects and organizes data about user-defined identifiers from PL/SQL source code. Because PL/Scope is a compiler-driven tool, you use it through interactive development environments (such as SQL Developer and JDeveloper), rather than directly.

PL/Scope enables the development of powerful and effective PL/Scope source code browsers that increase PL/SQL developer productivity by minimizing time spent browsing and understanding source code.

For more information about PL/Scope, see Chapter 7, "Using PL/Scope."

PL/SQL Hierarchical Profiler

The PL/SQL hierarchical profiler reports the dynamic execution profile of your PL/SQL program, organized by subprogram calls. It accounts for SQL and PL/SQL execution times separately. Each subprogram-level summary in the dynamic execution profile includes information such as number of calls to the subprogram, time spent in the subprogram itself, time spent in the subprogram's subtree (that is, in its descendent subprograms), and detailed parent-children information.

You can browse the generated HTML reports in any browser. The browser's navigational capabilities, combined with well chosen links, provide a powerful way to analyze performance of large applications, improve application performance, and lower development costs.

For a detailed description of PL/SQL hierarchical profiler, see Chapter 8, "Using the PL/SQL Hierarchical Profiler."

Oracle JDeveloper

Recent releases of Oracle JDeveloper have extensive features for debugging PL/SQL, Java, and multi-language programs. You can get Oracle JDeveloper as part of various Oracle product suites. Often, a more recent release is available as a download at `http://www.oracle.com/technology/`.

DBMS_OUTPUT Package

You can also debug stored subprograms and triggers using the Oracle package `DBMS_OUTPUT`. Put `PUT` and `PUT_LINE` statements in your code to output the value of variables and expressions to your terminal.

Privileges for Debugging PL/SQL and Java Stored Subprograms

Starting with Oracle Database 10*g*, a new privilege model applies to debugging PL/SQL and Java code running within the database. This model applies whether you are using Oracle JDeveloper, Oracle Developer, or any of the various third-party

PL/SQL or Java development environments, and it affects both the DBMS_DEBUG and DBMS_DEBUG_JDWP APIs.

For a session to connect to a debugger, the effective user at the time of the connect operation must have the DEBUG CONNECT SESSION system privilege. This effective user might be the owner of a DR subprogram involved in making the connect call.

When a debugger becomes connected to a session, the session login user and the enabled session-level roles are fixed as the privilege environment for that debugging connection. Any DEBUG or EXECUTE privileges needed for debugging must be granted to that combination of user and roles.

- To be able to display and change Java public variables or variables declared in a PL/SQL package specification, the debugging connection must be granted either EXECUTE or DEBUG privilege on the relevant code.
- To be able to either display and change private variables or breakpoint and run code lines step by step, the debugging connection must be granted DEBUG privilege on the relevant code

> **Caution:** The DEBUG privilege allows a debugging session to do anything that the subprogram being debugged could have done if that action had been included in its code.

In addition to these privilege requirements, the ability to stop on individual code lines and debugger access to variables are allowed only in code compiled with debug information generated. Use the PL/SQL compilation parameter PLSQL_DEBUG and the DEBUG keyword on statements such as ALTER PACKAGE to control whether the PL/SQL compiler includes debug information in its results. If not, variables are not accessible, and neither stepping nor breakpoints stop on code lines. The PL/SQL compiler never generates debug information for code hidden with the PL/SQL wrap utility.

> **See Also:** *Oracle Database PL/SQL Language Reference*, for information about the wrap utility

The DEBUG ANY PROCEDURE system privilege is equivalent to the DEBUG privilege granted on all objects in the database. Objects owned by SYS are included if the value of the O7_DICTIONARY_ACCESSIBILITY parameter is TRUE.

A debug role mechanism is available to carry privileges needed for debugging that are not normally enabled in the session. See the documentation on the DBMS_DEBUG and DBMS_DEBUG_JDWP packages for details on how to specify a debug role and any necessary related password.

The JAVADEBUGPRIV role carries the DEBUG CONNECT SESSION and DEBUG ANY PROCEDURE privileges. Grant it only with the care those privileges warrant.

> **Caution:** Granting DEBUG ANY PROCEDURE privilege, or granting DEBUG privilege on any object owned by SYS, means granting complete rights to the database.

Writing Low-Level Debugging Code

If you are writing code for part of a debugger, you might need to use packages such as DBMS_DEBUG_JDWP or DBMS_DEBUG.

DBMS_DEBUG_JDWP Package

The DBMS_DEBUG_JDWP package, provided starting with Oracle Database 9*g* Release 2 (9.2), provides a framework for multi-language debugging that is expected to supersede the DBMS_DEBUG package over time. It is especially useful for programs that combine PL/SQL and Java.

DBMS_DEBUG Package

The DBMS_DEBUG package, provided starting with Oracle8*i*, implements server-side debuggers and provides a way to debug server-side PL/SQL units. Several of the debuggers available, such as Oracle Procedure Builder and various third-party vendor solutions, use this API.

See Also:

- *Oracle Procedure Builder Developer's Guide*
- *Oracle Database PL/SQL Packages and Types Reference* for more information about the DBMS_DEBUG package and associated privileges
- *Oracle Database PL/SQL Packages and Types Reference* for more information about the DBMS_OUTPUT package and associated privileges
- The Oracle JDeveloper documentation for information about using package DBMS_DEBUG_JDWP
- *Oracle Database SQL Language Reference* for more details on privileges
- The PL/SQL page at http://www.oracle.com/technology/ for information about writing low-level debug code

Invoking Stored Subprograms

Stored PL/SQL subprograms can be invoked from many different environments. For example:

- Interactively, using an Oracle Database tool
- From the body of another subprogram
- From within an application (such as a SQL*Forms or a precompiler)
- From the body of a trigger

Stored PL/SQL functions (but not procedures) can also be invoked from within SQL statements. For details, see "Invoking Stored PL/SQL Functions from SQL Statements" on page 6-35.

Topics:

- Privileges Required to Invoke a Subprogram
- Invoking a Subprogram Interactively from Oracle Tools
- Invoking a Subprogram from Another Subprogram
- Invoking a Subprogram from a 3GL Application

See Also:
- *Oracle Database PL/SQL Language Reference* for information about invoking PL/SQL subprograms, including passing parameters.
- *Oracle Database PL/SQL Language Reference* for information about coding the body of a trigger

Privileges Required to Invoke a Subprogram

You do not need privileges to invoke:

- Standalone subprograms that you own
- Subprograms in packages that you own
- Public standalone subprograms
- Subprograms in public packages

To invoke a standalone or package subprogram owned by another user:

- You must have the `EXECUTE` privilege for the standalone subprogram or for the package containing the subprogram, or you must have the `EXECUTE ANY PROCEDURE` system privilege.
- When running a remote subprogram, you must be granted the `EXECUTE` privilege or `EXECUTE ANY PROCEDURE` system privilege directly, not through a role.
- You must include the name of the owner in the invocation. For example:

```
EXECUTE jdoe.Fire_emp (1043);
EXECUTE jdoe.Hire_fire.Fire_emp (1043);
```

- If the subprogram is a **definer's-rights (DR) subprogram**, then it runs with the privileges of the owner. The owner must have all the necessary object privileges for any referenced objects.
- If the subprogram is an **invoker's-rights (IR) subprogram**, then it runs with your privileges. You must have all the necessary object privileges for any referenced objects; that is, all objects accessed by the subprogram through external references that are resolved in your schema. You can hold these privileges either directly or through a role. Roles are enabled unless an IR subprogram is invoked directly or indirectly by a DR subprogram.

Invoking a Subprogram Interactively from Oracle Tools

You can invoke a subprogram interactively from an Oracle Database tool, such as SQL*Plus. Example 6–15 uses SQL*Plus to create a procedure and then invokes it in two different ways.

Example 6–15 Invoking a Subprogram Interactively with SQL*Plus

```
CREATE OR REPLACE PROCEDURE salary_raise (
  employee  EMPLOYEES.EMPLOYEE_ID%TYPE,
  increase  EMPLOYEES.SALARY%TYPE
)
IS
BEGIN
  UPDATE EMPLOYEES
  SET SALARY = SALARY + increase
  WHERE EMPLOYEE_ID = employee;
END;
```

```
/
```

Invoke procedure from within PL/SQL block:

```
BEGIN
  salary_raise(205, 200);
END;
/
```

Result:

```
PL/SQL procedure successfully completed.
```

Invoke procedure with EXECUTE statement:

```
EXECUTE salary_raise(205, 200);
```

Result:

```
PL/SQL procedure successfully completed.
```

Some interactive tools allow you to create session variables, which you can use for the duration of the session. Using SQL*Plus, Example 6–16 creates, uses, and prints a session variable.

Example 6–16 Creating and Using a Session Variable with SQL*Plus

```
-- Create function for later use:

CREATE OR REPLACE FUNCTION get_job_id (
  emp_id  EMPLOYEES.EMPLOYEE_ID%TYPE
) RETURN EMPLOYEES.JOB_ID%TYPE
IS
  job_id  EMPLOYEES.JOB_ID%TYPE;
BEGIN
  SELECT JOB_ID INTO job_id
  FROM EMPLOYEES
  WHERE EMPLOYEE_ID = emp_id;

  RETURN job_id;
END;
/
-- Create session variable:

VARIABLE job VARCHAR2(10);

-- Run function and store returned value in session variable:

EXECUTE :job := get_job_id(204);

PL/SQL procedure successfully completed.
```

SQL*Plus command:

```
PRINT job;
```

Result:

```
JOB
--------------------------------
PR_REP
```

See Also:

- *SQL*Plus User's Guide and Reference* for information about the EXECUTE command
- Your tools documentation for information about performing similar operations using your development tool

Invoking a Subprogram from Another Subprogram

A subprogram or a trigger can invoke another stored subprogram. In Example 6–17, the procedure print_mgr_name invokes the procedure print_emp_name.

Recursive subprogram invocations are allowed (that is, a subprogram can invoke itself).

Example 6–17 Invoking a Subprogram from Within Another Subprogram

```
-- Create procedure that takes employee's ID and prints employee's name:

CREATE OR REPLACE PROCEDURE print_emp_name (
  emp_id  EMPLOYEES.EMPLOYEE_ID%TYPE
)
IS
  fname  EMPLOYEES.FIRST_NAME%TYPE;
  lname  EMPLOYEES.LAST_NAME%TYPE;
BEGIN
  SELECT FIRST_NAME, LAST_NAME
  INTO fname, lname
  FROM EMPLOYEES
  WHERE EMPLOYEE_ID = emp_id;

  DBMS_OUTPUT.PUT_LINE (
    'Employee #' || emp_id || ':  ' || fname || ' ' || lname
  );
END;
/

-- Create procedure that takes employee's ID and prints manager's name:

CREATE OR REPLACE PROCEDURE print_mgr_name (
  emp_id  EMPLOYEES.EMPLOYEE_ID%TYPE
)
IS
  mgr_id  EMPLOYEES.MANAGER_ID%TYPE;
BEGIN
  SELECT MANAGER_ID
  INTO mgr_id
  FROM EMPLOYEES
  WHERE EMPLOYEE_ID = emp_id;

 DBMS_OUTPUT.PUT_LINE (
   'Manager of employee #' || emp_id || ' is:  '
 );

 print_emp_name(mgr_id);
END;
/
```

Invoke procedures:

```
BEGIN
  print_emp_name(200);
  print_mgr_name(200);
END;
/
```

Result:

```
Employee #200:  Jennifer Whalen
Manager of employee #200 is:
Employee #101:  Neena Kochhar
```

Invoking a Subprogram from a 3GL Application

A 3GL database application, such as a precompiler or an OCI application, can invoke a subprogram from within its own code.

Assume that the procedure `Fire_emp1` was created as follows:

```
CREATE OR REPLACE PROCEDURE fire_emp1 (Emp_id NUMBER) AS
  BEGIN
    DELETE FROM Emp_tab WHERE Empno = Emp_id;
  END;
```

To run a subprogram within the code of a precompiler application, you must use the `EXEC` call interface. For example, this statement invokes the `Fire_emp` procedure in the code of a precompiler application:

```
EXEC SQL EXECUTE
  BEGIN
    Fire_emp1(:Empnum);
  END;
END-EXEC;
```

> **See Also:** *Oracle Call Interface Programmer's Guide* for information about invoking PL/SQL subprograms from within 3GL applications

Invoking Remote Subprograms

Remote subprograms (standalone and package) can be invoked from within a subprogram, OCI application, or precompiler by specifying the remote subprogram name, a database link, and the parameters for the remote subprogram.

For example, this SQL*Plus statement invokes the procedure `fire_emp1`, which is located in the database and referenced by the local database link named `boston_server`:

```
EXECUTE fire_emp1@boston_server(1043);
```

You must specify values for all remote subprogram parameters, even if there are defaults. You cannot access remote package variables and constants.

Caution:

- Remote subprogram invocations use run-time binding. The user account to which you connect depends on the database link. (Stored subprograms use compile-time binding.)
- If a local subprogram invokes a remote subprogram, and a time stamp mismatch is found during execution of the local subprogram, then the remote subprogram is not run, and the local subprogram is invalidated.

Topics:

- Synonyms for Remote Subprograms
- Committing Transactions

See Also: "Handling Errors in Remote Subprograms" on page 6-26 for information about exception handling when invoking remote subprograms

Synonyms for Remote Subprograms

You can create a synonym for a remote subprogram name and database link, and then use the synonym to invoke the subprogram. For example:

```
CREATE SYNONYM synonym1 for fire_emp1@boston_server;

EXECUTE synonym1(1043);
/
```

The synonym enables you to invoke the remote subprogram from an Oracle Database tool application, such as a SQL*Forms application, as well from within a subprogram, OCI application, or precompiler.

Synonyms provide both data independence and location transparency. Synonyms permit applications to function without modification regardless of which user owns the object and regardless of which database holds the object. However, synonyms are not a substitute for privileges on database objects. Appropriate privileges must be granted to a user before the user can use the synonym.

Because subprograms defined within a package are not individual objects (the package is the object), synonyms cannot be created for individual subprograms within a package.

If you do not want to use a synonym, you can create a local subprogram to invoke the remote subprogram. For example:

```
CREATE OR REPLACE PROCEDURE local_procedure
  (arg IN NUMBER)
AS
BEGIN
  fire_emp1@boston_server(arg);
END;
/
DECLARE
  arg NUMBER;
BEGIN
  local_procedure(arg);
END;
```

/

See Also:

- *Oracle Database Concepts* for general information about synonyms
- *Oracle Database SQL Language Reference* for information about the `CREATE SYNONYM` statement

Committing Transactions

All invocations to remotely stored subprograms are assumed to perform updates; therefore, this type of referencing always requires two-phase commit of that transaction (even if the remote subprogram is read-only). Furthermore, if a transaction that includes a remote subprogram invocation is rolled back, then the work done by the remote subprogram is also rolled back.

A subprogram invoked remotely can usually run a `COMMIT`, `ROLLBACK`, or `SAVEPOINT` statement, the same as a local subprogram. However, there are some differences in action:

- If the transaction was originated by a database that is not an Oracle database, as might be the case in XA applications, these operations are not allowed in the remote subprogram.
- After doing one of these operations, the remote subprogram cannot start any distributed transactions of its own.
- If the remote subprogram does not commit or roll back its work, the commit is done implicitly when the database link is closed. In the meantime, further invocations to the remote subprogram are not allowed because it is still considered to be performing a transaction.

A **distributed transaction** modifies data on two or more databases. A distributed transaction is possible using a subprogram that includes two or more remote updates that access data on different databases. Statements in the construct are sent to the remote databases, and the execution of the construct succeeds or fails as a unit. If part of a distributed update fails and part succeeds, then a rollback (of the entire transaction or to a savepoint) is required to proceed. Consider this when creating subprograms that perform distributed updates.

Invoking Stored PL/SQL Functions from SQL Statements

Caution: Because SQL is a declarative language, rather than an imperative (or procedural) one, you cannot know how many times a function invoked from a SQL statement will run—even if the function is written in PL/SQL, an imperative language.

If your application requires that a function be executed a certain number of times, do not invoke that function from a SQL statement. Use a cursor instead.

For example, if your application requires that a function be called for each selected row, then open a cursor, select rows from the cursor, and call the function for each row. This guarantees that the number of calls to the function is the number of rows fetched from the cursor.

To be invoked from a SQL statement, a stored PL/SQL function must be declared either at schema level or in a package specification.

These SQL statements can invoke stored PL/SQL functions:

- `INSERT`
- `UPDATE`
- `DELETE`
- `SELECT`
- `CALL`

  (`CALL` can also invoke a stored PL/SQL procedure.)

To invoke a PL/SQL subprogram from SQL, you must either own or have `EXECUTE` privileges on the subprogram. To select from a view defined with a PL/SQL function, you must have `SELECT` privileges on the view. No separate `EXECUTE` privileges are necessary to select from the view.

For general information about invoking subprograms, including passing parameters, see *Oracle Database PL/SQL Language Reference*.

Topics:

- Why Invoke Stored PL/SQL Subprograms from SQL Statements?
- Where PL/SQL Functions Can Appear in SQL Statements
- When PL/SQL Functions Can Appear in SQL Expressions
- Controlling Side Effects

Why Invoke Stored PL/SQL Subprograms from SQL Statements?

Invoking PL/SQL subprograms in SQL statements can:

- Increase user productivity by extending SQL

  Expressiveness of the SQL statement increases where activities are too complex, too awkward, or unavailable with SQL.

- Increase query efficiency

  Functions used in the `WHERE` clause of a query can filter data using criteria that must otherwise be evaluated by the application.

- Manipulate character strings to represent special data types (for example, latitude, longitude, or temperature).
- Provide parallel query execution

  If the query is parallelized, then SQL statements in your PL/SQL subprogram might also be run in parallel (using the parallel query option).

Where PL/SQL Functions Can Appear in SQL Statements

A PL/SQL function can appear in a SQL statement wherever a built-in SQL function or an expression can appear in a SQL statement. For example:

- Select list of the `SELECT` statement
- Condition of the `WHERE` or `HAVING` clause
- `CONNECT BY`, `START WITH`, `ORDER BY`, or `GROUP BY` clause

- VALUES clause of the INSERT statement
- SET clause of the UPDATE statement

A PL/SQL table function (which returns a collection of rows) can appear in a SELECT statement instead of:

- Column name in the SELECT list
- Table name in the FROM clause

A PL/SQL function cannot appear in these contexts, which require unchanging definitions:

- CHECK constraint clause of a CREATE or ALTER TABLE statement
- Default value specification for a column

When PL/SQL Functions Can Appear in SQL Expressions

To be invoked from a SQL expression, a PL/SQL function must satisfy these requirements:

- It must be a row function, not a column (group) function; that is, its argument cannot be an entire column.
- Its formal parameters must be IN parameters, not OUT or IN OUT parameters.
- Its formal parameters and its return value (if any) must have Oracle built-in data types (such as CHAR, DATE, or NUMBER), not PL/SQL data types (such as BOOLEAN, RECORD, or TABLE).

  There is an exception to this rule: A formal parameter can have a PL/SQL data type if the corresponding actual parameter is implicitly converted to the data type of the formal parameter (as in Example 6–19).

The function in Example 6–18 satisfies the preceding requirements.

Example 6–18 PL/SQL Function in SQL Expression (Follows Rules)

```
DROP TABLE payroll;  -- in case it exists
CREATE TABLE payroll (
  srate  NUMBER,
  orate  NUMBER,
  acctno NUMBER
);

CREATE OR REPLACE FUNCTION gross_pay (
  emp_id  IN NUMBER,
  st_hrs  IN NUMBER := 40,
  ot_hrs  IN NUMBER := 0
) RETURN NUMBER
IS
  st_rate  NUMBER;
  ot_rate  NUMBER;
BEGIN
  SELECT srate, orate
  INTO st_rate, ot_rate
  FROM payroll
  WHERE acctno = emp_id;

 RETURN st_hrs * st_rate + ot_hrs * ot_rate;
END gross_pay;
/
```

In Example 6–19, the SQL statement CALL invokes the PL/SQL function f1, whose formal parameter and return value have PL/SQL data type PLS_INTEGER. The CALL statement succeeds because the actual parameter, 2, is implicitly converted to the data type PLS_INTEGER. If the actual parameter had a value outside the range of PLS_INTEGER, the CALL statement would fail.

Example 6–19 PL/SQL Function in SQL Expression (Exception to Rule)

```
CREATE OR REPLACE FUNCTION f1 (
  b IN PLS_INTEGER
) RETURN PLS_INTEGER
IS
BEGIN
  RETURN
    CASE
      WHEN b > 0  THEN  1
      WHEN b <= 0 THEN -1
      ELSE NULL
    END;
END f1;
/

VARIABLE x NUMBER;
CALL f1(b=>2) INTO :x;
PRINT x;
```

Result:

```
         X
----------
         1
```

Controlling Side Effects

The **purity** of a stored subprogram refers to the side effects of that subprogram on database tables or package variables. Side effects can prevent the parallelization of a query, yield order-dependent (and therefore, indeterminate) results, or require that package state be maintained across user sessions. Various side effects are not allowed when a function is invoked from a SQL query or DML statement.

Before Oracle Database 8*g* Release 1 (8.1), Oracle Database leveraged the PL/SQL compiler to enforce restrictions during the compilation of a stored subprogram or a SQL statement. As of Oracle Database 8*g* Release 1 (8.1), the compile-time restrictions were relaxed, and a smaller set of restrictions are enforced during execution.

This change provides uniform support for stored subprograms written in PL/SQL, Java, and C, and it allows programmers the most flexibility possible.

Topics:

- Restrictions
- Declaring a Function
- Parallel Query and Parallel DML
- PRAGMA RESTRICT_REFERENCES for Backward Compatibility

Restrictions

When a new SQL statement is run, checks are made to see if it is logically embedded within the execution of a running SQL statement. This occurs if the statement is run from a trigger or from a subprogram that was in turn invoked from the running SQL statement. In these cases, further checks determine if the new SQL statement is safe in the specific context.

These restrictions are enforced on subprograms:

- A subprogram invoked from a query or DML statement might not end the current transaction, create or rollback to a savepoint, or `ALTER` the system or session.
- A subprogram invoked from a query (`SELECT`) statement or from a parallelized DML statement might not run a DML statement or otherwise modify the database.
- A subprogram invoked from a DML statement might not read or modify the particular table being modified by that DML statement.

These restrictions apply regardless of what mechanism is used to run the SQL statement inside the subprogram or trigger. For example:

- They apply to a SQL statement invoked from PL/SQL, whether embedded directly in a subprogram or trigger body, run using the native dynamic mechanism (`EXECUTE IMMEDIATE`), or run using the `DBMS_SQL` package.
- They apply to statements embedded in Java with SQLJ syntax or run using JDBC.
- They apply to statements run with OCI using the callback context from within an "external" C function.

You can avoid these restrictions if the execution of the new SQL statement is not logically embedded in the context of the running statement. PL/SQL autonomous transactions provide one escape (see "Autonomous Transactions" on page 1-30). Another escape is available using OCI from an external C function, if you create a new connection rather than using the handle available from the `OCIExtProcContext` argument.

Declaring a Function

You can use the keywords `DETERMINISTIC` and `PARALLEL_ENABLE` in the syntax for declaring a function. These are optimization hints that inform the query optimizer and other software components about:

- Functions that need not be invoked redundantly
- Functions permitted within a parallelized query or parallelized DML statement

Only functions that are `DETERMINISTIC` are allowed in function-based indexes and in certain snapshots and materialized views.

A deterministic function depends solely on the values passed into it as arguments and does not reference or modify the contents of package variables or the database or have other side-effects. Such a function produces the same result value for any combination of argument values passed into it.

You place the `DETERMINISTIC` keyword after the return value type in a declaration of the function. For example:

```
CREATE OR REPLACE FUNCTION f1 (
  p1 NUMBER
) RETURN NUMBER DETERMINISTIC
IS
BEGIN
```

```
  RETURN p1 * 2;
END;
/
```

You might place this keyword in these places:

- On a function defined in a `CREATE FUNCTION` statement
- In a function declaration in a `CREATE PACKAGE` statement
- On a method declaration in a `CREATE TYPE` statement

Do not repeat the keyword on the function or method body in a `CREATE PACKAGE BODY` or `CREATE TYPE BODY` statement.

Certain performance optimizations occur on invocations of functions that are marked `DETERMINISTIC` without any other action being required. These features require that any function used with them be declared `DETERMINISTIC`:

- Any user-defined function used in a function-based index.
- Any function used in a materialized view, if that view is to qualify for Fast Refresh or is marked `ENABLE QUERY REWRITE`.

The preceding functions features attempt to use previously calculated results rather than invoking the function when it is possible to do so.

It is good programming practice to make functions that fall into these categories `DETERMINISTIC`:

- Functions used in a `WHERE`, `ORDER BY`, or `GROUP BY` clause
- Functions that `MAP` or `ORDER` methods of a SQL type
- Functions that help determine whether or where a row appears in a result set

Keep these points in mind when you create `DETERMINISTIC` functions:

- The database cannot recognize if the action of the function is indeed deterministic. If the `DETERMINISTIC` keyword is applied to a function whose action is not truly deterministic, then the result of queries involving that function is unpredictable.
- If you change the semantics of a `DETERMINISTIC` function and recompile it, then existing function-based indexes and materialized views report results for the prior version of the function. Thus, if you change the semantics of a function, you must manually rebuild any dependent function-based indexes and materialized views.

> **See Also:** *Oracle Database PL/SQL Language Reference* for `CREATE FUNCTION` restrictions

Parallel Query and Parallel DML

Oracle Database's parallel execution feature divides the work of running a SQL statement across multiple processes. Functions invoked from a SQL statement that is run in parallel might have a separate copy run in each of these processes, with each copy invoked for only the subset of rows that are handled by that process.

Each process has its own copy of package variables. When parallel execution begins, these are initialized based on the information in the package specification and body as if a user is logging into the system; the values in package variables are not copied from the original login session. And changes made to package variables are not automatically propagated between the various sessions or back to the original session. Java `STATIC` class attributes are similarly initialized and modified independently in each process. Because a function can use package (or Java `STATIC`) variables to

accumulate some value across the various rows it encounters, Oracle Database cannot assume that it is safe to parallelize the execution of all user-defined functions.

For SELECT statements in Oracle Database versions before 8.1.5, the parallel query optimization allowed functions noted as both RNPS and WNPS in a PRAGMA RESTRICT_REFERENCES declaration to run in parallel. Functions defined with CREATE FUNCTION statements had their code implicitly examined to determine if they were pure enough; parallelized execution might occur even though a pragma cannot be specified on these functions.

> **See Also:** "PRAGMA RESTRICT_REFERENCES for Backward Compatibility" on page 6-41

For DML statements in Oracle Database versions before 8.1.5, the parallelization optimization looked to see if a function was noted as having all four of RNDS, WNDS, RNPS and WNPS specified in a PRAGMA RESTRICT_REFERENCES declaration; those functions that were marked as neither reading nor writing to either the database or package variables could run in parallel. Again, those functions defined with a CREATE FUNCTION statement had their code implicitly examined to determine if they were actually pure enough; parallelized execution might occur even though a pragma cannot be specified on these functions.

Oracle Database versions 8.1.5 and later continue to parallelize those functions that earlier versions recognize as parallelizable. The PARALLEL_ENABLE keyword is the preferred way to mark your code as safe for parallel execution. This keyword is syntactically similar to DETERMINISTIC as described in "Declaring a Function" on page 6-39; it is placed after the return value type in a declaration of the function, as in:

```
CREATE OR REPLACE FUNCTION f1 (
  p1 NUMBER
) RETURN NUMBER PARALLEL_ENABLE
IS
BEGIN
  RETURN p1 * 2;
END;
/
```

A PL/SQL function defined with CREATE FUNCTION might still be run in parallel without any explicit declaration that it is safe to do so, if the system can determine that it neither reads nor writes package variables nor invokes any function that might do so. A Java method or C function is never seen by the system as safe to run in parallel, unless the programmer explicitly indicates PARALLEL_ENABLE on the call specification, or provides a PRAGMA RESTRICT_REFERENCES indicating that the function is sufficiently pure.

An additional run-time restriction is imposed on functions run in parallel as part of a parallelized DML statement. Such a function is not permitted to in turn run a DML statement; it is subject to the same restrictions that are enforced on functions that are run inside a query (SELECT) statement.

> **See Also:** Restrictions on page 6-39

PRAGMA RESTRICT_REFERENCES for Backward Compatibility

In Oracle Database versions before 8.1.5 (Oracle8*i*), programmers used PRAGMA RESTRICT_REFERENCES to assert the purity level of a subprogram. In subsequent versions, use the hints PARALLEL_ENABLE and DETERMINISTIC, instead, to communicate subprogram purity to Oracle Database.

You can remove PRAGMA RESTRICT_REFERENCES from your code. However, this pragma remains available for backward compatibility in these situations:

- When it is impossible or impractical to edit existing code to completely remove PRAGMA RESTRICT_REFERENCES.

  For example, if subprogram S1 depends on subprogram S2, and you do not remove the pragma from S1, then you might need the pragma in S2 to compile S1.

- When replacing PRAGMA RESTRICT_REFERENCES in existing code with hints PARALLEL_ENABLE and DETERMINISTIC would negatively affect the action of new, dependent code. (Use PRAGMA RESTRICT_REFERENCES to preserve the action of the existing code.)

An existing PL/SQL application can thus continue using the pragma even on new functionality, to ease integration with the existing code. Do not use the pragma in a new application.

If you use PRAGMA RESTRICT_REFERENCES, place it in a package specification, not in a package body. It must follow the declaration of a subprogram, but it need not follow immediately. Only one pragma can reference a given subprogram declaration.

To code the PRAGMA RESTRICT_REFERENCES, use this syntax:

```
PRAGMA RESTRICT_REFERENCES (
    Function_name, WNDS [, WNPS] [, RNDS] [, RNPS] [, TRUST] );
```

Where:

Option	Description
WNDS	The subprogram writes no database state (does not modify database tables).
RNDS	The subprogram reads no database state (does not query database tables).
WNPS	The subprogram writes no package state (does not change the values of package variables).
RNPS	The subprogram reads no package state (does not reference the values of package variables)
TRUST	The other restrictions listed in the pragma are not enforced; they are simply assumed to be true. This allows easy invocation from functions that have RESTRICT_REFERENCES declarations to those that do not.

You can pass the arguments in any order. If any SQL statement inside the subprogram body violates a rule, then you get an error when the statement is parsed.

In Example 6–20, the function compound_ neither reads nor writes database or package state; therefore, you can assert the maximum purity level. Always assert the highest purity level that a subprogram allows, so that the PL/SQL compiler never rejects the subprogram unnecessarily.

Example 6–20 PRAGMA RESTRICT_REFERENCES

```
DROP TABLE accounts; -- in case it exists
CREATE TABLE accounts (
  acctno   INTEGER,
  balance  NUMBER
);

INSERT INTO accounts (acctno, balance)
VALUES (12345, 1000.00);
```

```
CREATE OR REPLACE PACKAGE finance AS
  FUNCTION compound_ (
    years  IN NUMBER,
    amount IN NUMBER,
    rate   IN NUMBER
   ) RETURN NUMBER;
  PRAGMA RESTRICT_REFERENCES (compound_, WNDS, WNPS, RNDS, RNPS);
END finance;
/
CREATE PACKAGE BODY finance AS
  FUNCTION compound_ (
    years  IN NUMBER,
    amount IN NUMBER,
    rate   IN NUMBER
   ) RETURN NUMBER
   IS
   BEGIN
     RETURN amount * POWER((rate / 100) + 1, years);
   END compound_;
  -- No pragma in package body
END finance;
/
DECLARE
  interest NUMBER;
BEGIN
  SELECT finance.compound_(5, 1000, 6)
  INTO interest
  FROM accounts
  WHERE acctno = 12345;
END;
/
```

Topics:

- Using the Keyword TRUST
- Differences between Static and Dynamic SQL Statements
- Overloading Package Functions

Using the Keyword TRUST When PRAGMA RESTRICT REFERENCES includes the keyword TRUST, the restrictions listed in the pragma are assumed to be true, and not enforced.

When you invoke a subprogram that is in a section of code that does not use pragmas (such as a Java method), from a section of PL/SQL code that does use pragmas, specify PRAGMA RESTRICT REFERENCES with TRUST for either the invoked subprogram or the invoking subprogram.

In both Example 6–21 and Example 6–22, the PL/SQL function f invokes the Java procedure java_sleep. In Example 6–21, this is possible because java_sleep is declared to be WNDS with TRUST. In Example 6–22, it is possible because f is declared to be WNDS with TRUST, which allows it to invoke any subprogram.

Example 6–21 PRAGMA RESTRICT REFERENCES with TRUST on Invokee

```
CREATE OR REPLACE PACKAGE p IS
  PROCEDURE java_sleep (milli_seconds IN NUMBER)
  AS LANGUAGE JAVA NAME 'java.lang.Thread.sleep(long)';
  PRAGMA RESTRICT_REFERENCES(java_sleep,WNDS,TRUST);
```

```
  FUNCTION f (n NUMBER) RETURN NUMBER;
END p;
/
CREATE OR REPLACE PACKAGE BODY p IS
  FUNCTION f (
    n NUMBER
   ) RETURN NUMBER
   IS
   BEGIN
     java_sleep(n);
     RETURN n;
   END f;
END p;
/
```

Example 6–22 PRAGMA RESTRICT REFERENCES with TRUST on Invoker

```
CREATE OR REPLACE PACKAGE p IS
  PROCEDURE java_sleep (milli_seconds IN NUMBER)
  AS LANGUAGE JAVA NAME 'java.lang.Thread.sleep(long)';

  FUNCTION f (n NUMBER) RETURN NUMBER;
  PRAGMA RESTRICT_REFERENCES(f,WNDS,TRUST);
END p;
/
CREATE OR REPLACE PACKAGE BODY p IS
  FUNCTION f (
    n NUMBER
  ) RETURN NUMBER
  IS
  BEGIN
    java_sleep(n);
     RETURN n;
  END f;
END p;
/
```

Differences between Static and Dynamic SQL Statements Static `INSERT`, `UPDATE`, and `DELETE` statements do not violate `RNDS` if these statements do not explicitly read any database states, such as columns of a table. However, dynamic `INSERT`, `UPDATE`, and `DELETE` statements always violate `RNDS`, regardless of whether the statements explicitly read database states.

This `INSERT` statement violates `RNDS` if it is executed dynamically, but it does not violate `RNDS` if it is executed statically.

```
INSERT INTO my_table values(3, 'BOB');
```

This `UPDATE` statement always violates `RNDS` statically and dynamically, because it explicitly reads the column `name` of `my_table`.

```
UPDATE my_table SET id=777 WHERE name='BOB';
```

Overloading Package Functions PL/SQL lets you **overload** package (but not standalone) functions; that is, you can use the same name for different functions if their formal parameters differ in number, order, or data type family. However, `PRAGMA RESTRICT_REFERENCES` applies to only one function declaration (the most recently declared one).

In Example 6–23, the pragma applies to the second declaration of `valid`.

Example 6–23 Overloaded Package Function with PRAGMA RESTRICT_REFERENCES

```
CREATE OR REPLACE PACKAGE tests AS
  FUNCTION valid (x NUMBER) RETURN CHAR;
  FUNCTION valid (x DATE) RETURN CHAR;
  PRAGMA RESTRICT_REFERENCES (valid, WNDS);
END;
/
```

Returning Large Amounts of Data from a Function

In a data warehousing environment, you might use PL/SQL functions to transform large amounts of data. Perhaps the data is passed through a series of transformations, each performed by a different function. PL/SQL table functions let you perform such transformations without significant memory overhead or the need to store the data in tables between each transformation stage. These functions can accept and return multiple rows, can return rows as they are ready rather than all at once, and can be parallelized.

> **See Also:** *Oracle Database PL/SQL Language Reference* for more information about performing multiple transformations with pipelined table functions

Coding Your Own Aggregate Functions

To analyze a set of rows and compute a result value, you can code your own aggregate function that works the same as a built-in aggregate like `SUM`:

- Define an ADT that defines these member functions:
  - `ODCIAggregateInitialize`
  - `ODCIAggregateIterate`
  - `ODCIAggregateMerge`
  - `ODCIAggregateTerminate`
- Code the member functions. In particular, `ODCIAggregateIterate` accumulates the result as it is invoked for each row that is processed. Store any intermediate results using the attributes of the ADT.
- Create the aggregate function, and associate it with the ADT.
- Call the aggregate function from SQL queries, DML statements, or other places that you might use the built-in aggregates. You can include typical options such as `DISTINCT` and `ALL` in the invocation of the aggregate function.

> **See Also:** *Oracle Database Data Cartridge Developer's Guide* for more information about user-defined aggregate functions

7

Using PL/Scope

PL/Scope is a compiler-driven tool that collects data about identifiers in PL/SQL source code at program-unit compilation time and makes it available in static data dictionary views. The collected data includes information about identifier types, usages (declaration, definition, reference, call, assignment) and the location of each usage in the source code.

PL/Scope enables the development of powerful and effective PL/Scope source code browsers that increase PL/SQL developer productivity by minimizing time spent browsing and understanding source code.

PL/Scope is intended for application developers, and is usually used in the environment of a development database.

Note: PL/Scope cannot collect data for a PL/SQL unit whose source code is wrapped. For information about wrapping PL/SQL source code, see *Oracle Database PL/SQL Language Reference*.

Topics:

- Specifying Identifier Collection
- PL/Scope Identifier Data for STANDARD and DBMS_STANDARD
- How Much Space is PL/Scope Data Using?
- Viewing PL/Scope Data
- Identifier Types that PL/Scope Collects
- Usages that PL/Scope Reports
- Sample PL/Scope Session

Specifying Identifier Collection

By default, PL/Scope does not collect data for identifiers in the PL/SQL source program. To have PL/Scope collect data for all identifiers in the PL/SQL source program, including identifiers in package bodies, set the PL/SQL compilation parameter `PLSCOPE_SETTINGS` to `'IDENTIFIERS:ALL'`.

Note: Collecting all identifiers might generate large amounts of data and slow compile time.

PL/Scope stores the data that it collects in the SYSAUX tablespace. If the SYSAUX tablespace is unavailable, and you compile a program unit with PLSCOPE_SETTINGS='IDENTIFIERS:ALL', PL/Scope does not collect data for the compiled object. The compiler does not issue a warning, but it saves a warning in USER_ERRORS.

See Also:

- *Oracle Database Reference* for information about PLSCOPE_SETTINGS
- *Oracle Database PL/SQL Language Reference* for information about PL/SQL compilation parameters

PL/Scope Identifier Data for STANDARD and DBMS_STANDARD

The packages STANDARD and DBMS_STANDARD declare and define base types, such as VARCHAR2 and NUMBER, and subprograms such as RAISE_APPLICATION_ERROR. If your database has PL/Scope identifier data for these packages, PL/Scope can track your usage of the identifiers that these packages create.

Do You Need STANDARD and DBMS_STANDARD Identifier Data?

You can use PL/Scope without STANDARD and DBMS_STANDARD identifier data. You need this data only if you must know where your code uses the base types or subprograms that these packages create—for example, to know where your code uses the base type BINARY_INTEGER, so that you can substitute PLS_INTEGER.

Does Your Database Have STANDARD and DBMS_STANDARD Identifier Data?

A newly created Oracle 11.1.0.7 database, or a database that was upgraded to 11.1.0.7 from 10.2, has PL/Scope identifier data for the packages STANDARD and DBMS_STANDARD. A database that was upgraded to 11.1.0.7 from 11.1.0.6 does not have this data.

To see if your database has this data, use the query in Example 7–1.

Example 7–1 shows what the query returns when the database has PL/Scope identifier data for STANDARD and DBMS_STANDARD.

Example 7–1 Is STANDARD and DBMS_STANDARD PL/Scope Identifier Data Available?

Query:

```
SELECT UNIQUE OBJECT_NAME
FROM ALL_IDENTIFIERS
WHERE OBJECT_NAME IN ('STANDARD', 'DBMS_STANDARD')
AND OWNER='SYS'
ORDER BY OBJECT_NAME;
```

Result:

```
OBJECT_NAME
------------------------------
DBMS_STANDARD
STANDARD

2 rows selected.
```

If the query in Example 7–1 selects no rows, then the database does not have PL/Scope identifier data for the packages STANDARD and DBMS_STANDARD. To collect this data,

a DBA must recompile the packages STANDARD and DBMS_STANDARD, as explained in "Recompiling STANDARD and DBMS_STANDARD" on page 7-3.

Recompiling STANDARD and DBMS_STANDARD

A DBA can use this procedure to recompile the packages STANDARD and DBMS_STANDARD:

> **Note:** This procedure invalidates and revalidates (by recompiling) every PL/SQL object in the database.

1. Connect to the database, shut it down, and then start it in UPGRADE mode:

```
CONNECT / AS SYSDBA;
SHUTDOWN IMMEDIATE;
STARTUP PFILE=parameter_initialization_file UPGRADE;
```

2. Have PL/Scope collect data for all identifiers in the packages STANDARD and DBMS_STANDARD:

```
ALTER SESSION SET PLSCOPE_SETTINGS='IDENTIFIERS:ALL';
```

3. Invalidate and recompile the database:

```
@?/rdbms/admin/utlirp.sql
```

   Now all PL/SQL objects in the database are invalid except STANDARD and DBMS_STANDARD, which were recompiled with PLSCOPE_SETTINGS='IDENTIFIERS:ALL'.

4. (Optional) Invalidate any other PL/SQL objects to be recompiled with PLSCOPE_SETTINGS='IDENTIFIERS:ALL', using a script similar to this.

   Customize the query on lines 5 through 9 to invalidate only those objects for which you need PL/Scope identifier data. Collecting all identifiers for all objects, as this script does, might generate large amounts of data and slow compile time:

```
DECLARE
  TYPE ObjIDArray IS TABLE OF NUMBER INDEX BY BINARY_INTEGER;
  ObjIDs  ObjIDArray;
BEGIN
  SELECT object_id BULK COLLECT INTO ObjIDs
  FROM ALL_OBJECTS
  WHERE object_type IN
    (SELECT DISTINCT TYPE
     FROM ALL_PLSQL_OBJECT_SETTINGS);
  FOR i IN 1..SQL%ROWCOUNT LOOP
    BEGIN
      DBMS_UTILITY.INVALIDATE(ObjIDs(i),
        'PLSCOPE_SETTINGS=IDENTIFIERS:ALL REUSE SETTINGS');
      NULL;
    END;
  END LOOP;
END;
/
```

Notes: In the preceding script:

- Do not substitute `ObjIDs.LAST` for `SQL%ROWCOUNT`, because `ObjIDs` attributes are dependent on a package that is locked by the anonymous block.
- If your database is large, do not substitute a cursor `FOR LOOP` for the `BULK COLLECT` statement, or you will run out of resources.

5. Shut down the database, and then start it in `NORMAL` mode:

```
SHUTDOWN IMMEDIATE;
STARTUP PFILE=parameter_initialization_file;
```

6. For any remaining invalid PL/SQL objects, do either of these:
   - Allow them to be recompiled automatically, as they are referenced.

     (This can be slow if there are complex dependencies.)
   - Run the script `utlrp.sql` to recompile the invalid PL/SQL objects, as explained in "Running utlrp.sql to Recompile Invalid PL/SQL Objects" on page 7-4.

Running utlrp.sql to Recompile Invalid PL/SQL Objects

If the database was restarted in `NORMAL` mode (step 5 on page 7-4), then a DBA, or a user who has been granted the DBA role, can use this procedure:

1. Connect to the database as `SYS`:

```
CONNECT / AS SYS;
```

2. Run the script `utlrp.sql`:

```
@?/rdbms/admin/utlrp.sql
```

   If the script gives you any instructions, follow them, and then run the script again.

   If the script terminates abnormally without giving any instructions, run it again.

How Much Space is PL/Scope Data Using?

PL/Scope stores its data in the `SYSAUX` tablespace. If you are logged on as `SYSDBA`, you can use the query in Example 7–2 to display the amount of space that PL/Scope data is using.

Example 7–2 How Much Space is PL/Scope Data Using?

Query:

```
SELECT SPACE_USAGE_KBYTES
FROM V$SYSAUX_OCCUPANTS
WHERE OCCUPANT_NAME='PL/SCOPE';
```

Result:

```
SPACE_USAGE_KBYTES
------------------
              1600

1 row selected.
```

For information about managing the SYSAUX tablespace, see *Oracle Database Administrator's Guide*.

Viewing PL/Scope Data

To view the data that PL/Scope has collected, you can use either:

- Static Data Dictionary Views
- Demo Tool
- SQL Developer

Static Data Dictionary Views

The static data dictionary views *_IDENTIFIERS display information about PL/Scope identifiers, including their types and usages. For general information about these views, see *Oracle Database Reference*.

Topics:

- Unique Keys
- Context
- Signature

Unique Keys

Each row of a *_IDENTIFIERS view represents a unique usage of an identifier in the PL/SQL unit. In each of these views, these are equivalent unique keys within a compilation unit:

- LINE, COL, and USAGE
- USAGE_ID

For the usages in the *_IDENTIFIERS views, see "Usages that PL/Scope Reports" on page 7-9.

Note: An identifier that is passed to a subprogram in IN OUT mode has two rows in *_IDENTIFIERS: a REFERENCE usage (corresponding to IN) and an ASSIGNMENT usage (corresponding to OUT).

Context

Context is useful for discovering relationships between usages. Except for top-level schema object declarations and definitions, every usage of an identifier happens within the context of another usage. For example:

- A local variable declaration happens within the context of a top-level procedure declaration.
- If an identifier is declared as a variable, such as x VARCHAR2(10), the USAGE_CONTEXT_ID of the VARCHAR2 type reference contains the USAGE_ID of the x declaration, allowing you to associate the variable declaration with its type.

In other words, USAGE_CONTEXT_ID is a reflexive foreign key to USAGE_ID, as Example 7–3 shows.

Example 7–3 USAGE_CONTEXT_ID and USAGE_ID

```
ALTER SESSION SET PLSCOPE_SETTINGS = 'IDENTIFIERS:ALL';

CREATE OR REPLACE PROCEDURE a (p1 IN BOOLEAN) IS
  v PLS_INTEGER;
BEGIN
  v := 42;
  DBMS_OUTPUT.PUT_LINE(v);
  RAISE_APPLICATION_ERROR (-20000, 'Bad');
EXCEPTION
  WHEN Program_Error THEN NULL;
END a;
/
CREATE OR REPLACE PROCEDURE b (p2 OUT PLS_INTEGER, p3 IN OUT VARCHAR2) IS
  n NUMBER;
  q BOOLEAN := TRUE;
BEGIN
  FOR j IN 1..5 LOOP
    a(q); a(TRUE); a(TRUE);
    IF j > 2 THEN
       GOTO z;
    END IF;
  END LOOP;
<<z>> DECLARE
  d CONSTANT CHAR(1) := 'X';
  BEGIN
    SELECT COUNT(*) INTO n FROM Dual WHERE Dummy = d;
  END z;
END b;
/
WITH v AS (
  SELECT    Line,
            Col,
            INITCAP(NAME) Name,
            LOWER(TYPE)   Type,
            LOWER(USAGE)  Usage,
            USAGE_ID,
            USAGE_CONTEXT_ID
    FROM USER_IDENTIFIERS
      WHERE Object_Name = 'B'
        AND Object_Type = 'PROCEDURE'
)
SELECT RPAD(LPAD(' ', 2*(Level-1)) ||
                Name, 20, '.')||' '||
                RPAD(Type, 20)||
                RPAD(Usage, 20)
                IDENTIFIER_USAGE_CONTEXTS
  FROM v
  START WITH USAGE_CONTEXT_ID = 0
  CONNECT BY PRIOR USAGE_ID = USAGE_CONTEXT_ID
  ORDER SIBLINGS BY Line, Col
/
```

Result:

```
IDENTIFIER_USAGE_CONTEXTS
-------------------------------------------------------------
B................... procedure           declaration
  B................. procedure           definition
    P2............... formal out          declaration
```

```
    Pls_Integer... subtype             reference
  P3............. formal in out       declaration
    Varchar2...... character datatype  reference
  N............... variable            declaration
    Number........ number datatype     reference
  Q............... variable            declaration
    Q............. variable            assignment
    Boolean....... boolean datatype    reference
  J............... iterator            declaration
    A............. procedure           call
      Q........... variable            reference
    A............. procedure           call
    A............. procedure           call
    J............. iterator            reference
    Z............. label               reference
  Z............... label               declaration
    D............. constant            declaration
      D........... constant            assignment
      Char........ subtype             reference
    N............. variable            assignment
    D............. constant            reference

24 rows selected.
```

Signature

The signature of an identifier is unique, within and across program units. That is, the signature distinguishes the identifier from other identifiers with the same name, whether they are defined in the same program unit or different program units.

For the program unit in Example 7–4, which has two identifiers named `p`, the static data dictionary view `USER_IDENTIFIERS` has several rows in which `NAME` is `p`, but in these rows, `SIGNATURE` varies. The rows associated with the outer procedure `p` have one signature, and the rows associated with the inner procedure `p` have another signature. If program unit `q` calls procedure `p`, the `USER_IDENTIFIERS` view for `q` has a row in which `NAME` is `p` and `SIGNATURE` is the signature of the outer procedure `p`.

Example 7–4 Program Unit with Two Identifiers Named p

```
CREATE OR REPLACE PROCEDURE p IS
  PROCEDURE p IS
  BEGIN
    DBMS_OUTPUT.PUT_LINE('Inner p');
  END p;
BEGIN
  DBMS_OUTPUT.PUT_LINE('Outer p');
  p();
END p;
/
```

Demo Tool

`$ORACLE_HOME/plsql/demo/plscopedemo.sql` is an HTML-based demo implemented as a PL/SQL Web Application using the PL/SQL Web Toolkit. For more information about PL/SQL Web Applications, see "Implementing PL/SQL Web Applications" on page 9-2.

SQL Developer

PL/Scope is a feature of SQL Developer. For information about using PL/Scope from SQL Developer, see the SQL Developer online documentation.

Identifier Types that PL/Scope Collects

Table 7–1 shows the identifier types that PL/Scope collects, in alphabetical order. The identifier types in Table 7–1 appear in the `TYPE` column of the `*_IDENTIFIER` static data dictionary views, which are described in *Oracle Database Reference*.

> **Note:** Identifiers declared in compilation units that were not compiled with `PLSCOPE_SETTINGS='IDENTIFIERS:ALL'` do not appear in `*_IDENTIFIER` static data dictionary views.

Table 7–1 Identifier Types that PL/Scope Collects

TYPE Column Value	Comment
ASSOCIATIVE ARRAY	
CONSTANT	
CURSOR	
BFILE DATATYPEBLOB DATATYPEBOOLEAN DATATYPECHARACTER DATATYPECLOB DATATYPEDATE DATATYPEINTERVAL DATATYPENUMBER DATATYPETIME DATATYPETIMESTAMP DATATYPE	Each DATATYPE is a base type declared in package STANDARD.
EXCEPTION	
FORMAL INFORMAL IN OUTFORMAL OUT	
FUNCTION	
INDEX TABLE	
ITERATOR	An iterator is the index of a FOR loop.
LABEL	A label declaration also acts as a context.
LIBRARY	
NESTED TABLE	
OBJECT	
OPAQUE	Examples of internal opaque types are ANYDATA and XMLType.
PACKAGE	
PROCEDURE	
RECORD	
REFCURSOR	
SUBTYPE	

Table 7–1 (Cont.) Identifier Types that PL/Scope Collects

TYPE Column Value	Comment
SYNONYM	PL/Scope does not resolve the base object name of a synonym. To find the base object name of a synonym, query *_SYNONYMS.
TRIGGER	
UROWID	
VARRAY	
VARIABLE	Can be object attribute, local variable, package variable, or record field.

Usages that PL/Scope Reports

Table 7–2 shows the usages that PL/Scope reports, in alphabetical order. The identifier types in Table 7–2 appear in the USAGE column of the *_IDENTIFIER static data dictionary views, which are described in *Oracle Database Reference*.

Table 7–2 Usages that PL/Scope Reports

USAGE Column Value	Description
ASSIGNMENT	An assignment can be made only to an identifier that can have a value, such as a VARIABLE. Examples of assignments are: ▪ Using an identifier to the left of an assignment operator ▪ Using an identifier in the INTO clause of a FETCH statement ▪ Passing an identifier to a subprogram by reference (OUT mode) ▪ Using an identifier as the bind argument in the USING clause of an EXECUTE IMMEDIATE statement in either OUT or IN OUT mode An identifier that is passed to a subprogram in IN OUT mode has both a REFERENCE usage (corresponding to IN) and an ASSIGNMENT usage (corresponding to OUT).
CALL	In the context of PL/Scope, a CALL is an operation that pushes a call onto the call stack; that is: ▪ A call to a FUNCTION or PROCEDURE ▪ Running or fetching a cursor identifier (a logical call to SQL) A GOTO statement or raise of an exception is not a CALL, because neither pushes a call onto the call stack.
DECLARATION	A DECLARATION tells the compiler that an identifier exists, and each identifier has exactly one DECLARATION. Each DECLARATION can have an associated data type. For a loop index declaration, LINE and COL (in *_IDENTIFIERS views) are the line and column of the FOR clause that implicitly declares the loop index. For a label declaration, LINE and COL are the line and column on which the label appears (and is implicitly declared) within the delimiters << and >>.

Table 7–2 (Cont.) Usages that PL/Scope Reports

USAGE Column Value	Description
DEFINITION	A DEFINITION tells the compiler how to implement or use a previously declared identifier. Each of these types of identifiers has a DEFINITION: ■ EXCEPTION (can have multiple definitions) ■ FUNCTION ■ OBJECT ■ PACKAGE ■ PROCEDURE ■ TRIGGER For a top-level identifier only, the DEFINITION and DECLARATION are in the same place.
REFERENCE	A REFERENCE uses an identifier without changing its value. Examples of references are: ■ Raising an exception identifier ■ Using a type identifier in the declaration of a variable or formal parameter ■ Using a variable identifier whose type contains fields to access a field. For example, in myrecordvar.myfield := 1, a reference is made to myrecordvar, and an assignment is made to myfield. ■ Using a cursor for any purpose except FETCH ■ Passing an identifier to a subprogram by value (IN mode) ■ Using an identifier as the bind argument in the USING clause of an EXECUTE IMMEDIATE statement in either IN or IN OUT mode An identifier that is passed to a subprogram in IN OUT mode has both a REFERENCE usage (corresponding to IN) and an ASSIGNMENT usage (corresponding to OUT).

Sample PL/Scope Session

In this sample session, assume that you are logged in as HR.

1. Set the session parameter:

```
ALTER SESSION SET PLSCOPE_SETTINGS='IDENTIFIERS:ALL';
```

2. Create this package:

```
CREATE OR REPLACE PACKAGE PACK1 IS
  TYPE r1 is RECORD (rf1 VARCHAR2(10));
  FUNCTION F1(fp1 NUMBER) RETURN NUMBER;
  PROCEDURE P1(pp1 VARCHAR2);
END PACK1;
/
CREATE OR REPLACE PACKAGE BODY PACK1 IS
  FUNCTION F1(fp1 NUMBER) RETURN NUMBER IS
    a NUMBER := 10;
  BEGIN
    RETURN a;
  END F1;
  PROCEDURE P1(pp1 VARCHAR2) IS
    pr1 r1;
```

```
  BEGIN
    pr1.rf1 := pp1;
  END;
END PACK1;
/
```

3. Verify that PL/Scope collected all identifiers for the package body:

```
SELECT PLSCOPE_SETTINGS
FROM USER_PLSQL_OBJECT_SETTINGS
WHERE NAME='PACK1' AND TYPE='PACKAGE BODY'
```

Result:

```
PLSCOPE_SETTINGS
----------------------------------------------------------------------
IDENTIFIERS:ALL
```

4. Display unique identifiers in `HR` by querying for all `DECLARATION` usages. For example, to see all unique identifiers with name like `%1`, use these SQL*Plus formatting commands and this query:

```
COLUMN NAME FORMAT A6
COLUMN SIGNATURE FORMAT A32
COLUMN TYPE FORMAT A9

SELECT NAME, SIGNATURE, TYPE
FROM USER_IDENTIFIERS
WHERE NAME LIKE '%1' AND USAGE='DECLARATION'
ORDER BY OBJECT_TYPE, USAGE_ID;
```

Result is similar to:

```
NAME   SIGNATURE                         TYPE
------ -------------------------------- ---------
PACK1  41820FA4D5EF6BE707895178D0C5C4EF PACKAGE
R1     EEBB6849DEE31BC77BF186EBAE5D4E2D RECORD
RF1    41D70040337349634A7F547BC83517C7 VARIABLE
F1     4559CF050A5F5C3E5F5FFDD0D9D55EFA FUNCTION
FP1    CAC3474C112DBEC67AB926978D9A16C1 FORMAL IN
P1     B7C0576BA4D00C33A65CC0C64CADAB89 PROCEDURE
PP1    6B74CF95A5B7377A735925DFAA280266 FORMAL IN
FP1    98EB63B8A4AFEB5EF94D50A20165D6CC FORMAL IN
PP1    AD89FE0EAE9CE5D6D48AA4684E0D57DF FORMAL IN
PR1    1B5117F30E8DAE0261A02CAA5E33883F VARIABLE

10 rows selected.
```

The `*_IDENTIFIERS` static data dictionary views display only basic type names; for example, the `TYPE` of a local variable or record field is `VARIABLE`. To determine the exact type of a `VARIABLE`, you must use its `USAGE_CONTEXT_ID`.

5. Find all local variables:

```
COLUMN VARIABLE_NAME FORMAT A13
COLUMN CONTEXT_NAME FORMAT A12

SELECT a.NAME variable_name,
        b.NAME context_name,
        a.SIGNATURE
FROM USER_IDENTIFIERS a, USER_IDENTIFIERS b
WHERE a.USAGE_CONTEXT_ID = b.USAGE_ID
```

```
AND a.TYPE = 'VARIABLE'
AND a.USAGE = 'DECLARATION'
AND a.OBJECT_NAME = 'PACK1'
AND a.OBJECT_NAME = b.OBJECT_NAME
AND a.OBJECT_TYPE =  b.OBJECT_TYPE
AND (b.TYPE = 'FUNCTION' or b.TYPE = 'PROCEDURE')
ORDER BY a.OBJECT_TYPE, a.USAGE_ID;
```

Result is similar to:

```
VARIABLE_NAME CONTEXT_NAME SIGNATURE
------------- ------------ --------------------------------
A             F1           1691C6B3C951FCAA2CBEEB47F85CF128
PR1           P1           1B5117F30E8DAE0261A02CAA5E33883F

2 rows selected.
```

6. Find all usages performed on the local variable `A`:

```
COLUMN USAGE FORMAT A11
COLUMN USAGE_ID FORMAT 999
COLUMN OBJECT_NAME FORMAT A11
COLUMN OBJECT_TYPE FORMAT A12

SELECT USAGE, USAGE_ID, OBJECT_NAME, OBJECT_TYPE
FROM USER_IDENTIFIERS
WHERE SIGNATURE='1691C6B3C951FCAA2CBEEB47F85CF128'  -- signature of A
ORDER BY OBJECT_TYPE, USAGE_ID;
```

Result:

```
USAGE       USAGE_ID OBJECT_NAME OBJECT_TYPE
----------- -------- ----------- ------------
DECLARATION        6 PACK1       PACKAGE BODY
ASSIGNMENT         8 PACK1       PACKAGE BODY
REFERENCE          9 PACK1       PACKAGE BODY

3 rows selected.
```

The usages performed on the local identifier `A` are the identifier declaration (`USAGE_ID` 6), an assignment (`USAGE_ID` 8), and a reference (`USAGE_ID` 9).

7. From the declaration of the local identifier `A`, find its type:

```
COLUMN NAME FORMAT A6
COLUMN TYPE FORMAT A15

SELECT a.NAME, a.TYPE
FROM USER_IDENTIFIERS a, USER_IDENTIFIERS b
WHERE a.USAGE = 'REFERENCE'
AND a.USAGE_CONTEXT_ID = b.USAGE_ID
AND b.USAGE = 'DECLARATION'
AND b.SIGNATURE = '4559CF050A5F5C3E5F5FFDD0D9D55EFA'  -- signature of F1
AND a.OBJECT_TYPE = b.OBJECT_TYPE
AND a.OBJECT_NAME = b.OBJECT_NAME;
```

Result:

```
NAME   TYPE
------ ---------------
NUMBER NUMBER DATATYPE
```

```
1 row selected.
```

> **Note:** This query produces the output shown only if your database has PL/Scope identifier data for the packages STANDARD and DBMS_STANDARD. For more information, see "PL/Scope Identifier Data for STANDARD and DBMS_STANDARD" on page 7-2.

8. Find out where the assignment to local identifier A occurred:

```
SELECT LINE, COL, OBJECT_NAME, OBJECT_TYPE
FROM USER_IDENTIFIERS
WHERE SIGNATURE='1691C6B3C951FCAA2CBEEB47F85CF128'  -- signature of A
AND USAGE='ASSIGNMENT';
```

Result:

```
      LINE        COL OBJECT_NAME OBJECT_TYPE
---------- ---------- ----------- ------------
         3          5 PACK1       PACKAGE BODY

1 row selected.
```

8

Using the PL/SQL Hierarchical Profiler

You can use the PL/SQL hierarchical profiler to identify bottlenecks and performance-tuning opportunities in PL/SQL applications.

The profiler reports the dynamic execution profile of a PL/SQL program organized by function calls, and accounts for SQL and PL/SQL execution times separately. No special source or compile-time preparation is required; any PL/SQL program can be profiled.

This chapter describes the PL/SQL hierarchical profiler and explains how to use it to collect and analyze profile data for a PL/SQL program.

Topics:

- Overview of PL/SQL Hierarchical Profiler
- Collecting Profile Data
- Understanding Raw Profiler Output
- Analyzing Profile Data
- plshprof Utility

Overview of PL/SQL Hierarchical Profiler

Nonhierarchical (**flat**) profilers record the time that a program spends within each subprogram—the **function time** or **self time** of each subprogram. Function time is helpful, but often inadequate. For example, it is helpful to know that a program spends 40% of its time in the subprogram `INSERT_ORDER`, but it is more helpful to know which subprograms call `INSERT_ORDER` often and the total time the program spends under `INSERT_ORDER` (including its descendant subprograms). Hierarchical profilers provide such information.

The PL/SQL hierarchical profiler:

- Reports the dynamic execution profile of your PL/SQL program, organized by subprogram calls
- Accounts for SQL and PL/SQL execution times separately
- Requires no special source or compile-time preparation
- Stores results in database tables (**hierarchical profiler tables**) for custom report generation by integrated development environment (IDE) tools (such as SQL Developer and third-party tools)
- Provides subprogram-level execution summary information, such as:

- Number of calls to the subprogram
- Time spent in the subprogram itself (**function time** or **self time**)
- Time spent in the subprogram itself and in its descendent subprograms (**subtree time**)
- Detailed parent-children information, for example:
  - All callers of a given subprogram (parents)
  - All subprograms that a given subprogram called (children)
  - How much time was spent in subprogram *x* when called from *y*
  - How many calls to subprogram *x* were from *y*

The PL/SQL hierarchical profiler is implemented by the `DBMS_HPROF` package and has two components:

- Data collection

  The data collection component is an intrinsic part of the PL/SQL Virtual Machine. The `DBMS_HPROF` package provides APIs to turn hierarchical profiling on and off and write the **raw profiler output** to a file.

- Analyzer

  The analyzer component processes the raw profiler output and stores the results in hierarchical profiler tables.

> **Note:** To generate simple HTML reports directly from raw profiler output, without using the Analyzer, you can use the `plshprof` command-line utility.

Collecting Profile Data

To collect profile data from your PL/SQL program for the PL/SQL hierarchical profiler, follow these steps:

1. Ensure that you have these privileges:
   - `EXECUTE` privilege on the `DBMS_HPROF` package
   - `WRITE` privilege on the directory that you specify when you call `DBMS_HPROF.START_PROFILING`
2. Use the `DBMS_HPROF.START_PROFILING` PL/SQL API to start hierarchical profiler data collection in a session.
3. Run your PL/SQL program long enough to get adequate code coverage.

   To get the most accurate measurements of elapsed time, avoid unrelated activity on the system on which your PL/SQL program is running.
4. Use the `DBMS_HPROF.STOP_PROFILING` PL/SQL API to stop hierarchical profiler data collection.

For more information about `DBMS_HPROF.START_PROFILING` and `DBMS_HPROF.STOP_PROFILING`, see *Oracle Database PL/SQL Packages and Types Reference*.

Consider this PL/SQL procedure, `test`:

```
CREATE OR REPLACE PROCEDURE test IS
  n NUMBER;
```

```
  PROCEDURE foo IS
  BEGIN
    SELECT COUNT(*) INTO n FROM EMPLOYEES;
  END foo;

BEGIN  -- test
  FOR i IN 1..3 LOOP
    foo;
  END LOOP;
END test;
/
```

The SQL script in Example 8–1 profiles the execution of the PL/SQL procedure `test`.

Example 8–1 Profiling a PL/SQL Procedure

```
BEGIN
  /* Start profiling.
     Write raw profiler output to file test.trc in a directory
     that is mapped to directory object PLSHPROF_DIR
     (see note after example). */

  DBMS_HPROF.START_PROFILING('PLSHPROF_DIR', 'test.trc');
END;
/
-- Run procedure to be profiled
BEGIN
  test;
END;
/
BEGIN
  -- Stop profiling
  DBMS_HPROF.STOP_PROFILING;
END;
/
```

Note: A directory object is an alias for a file system path name. For example, if you are connected to the database `AS SYSDBA`, this `CREATE DIRECTORY` statement creates the directory object `PLSHPROF_DIR` and maps it to the file system directory `/private/plshprof/results`:

```
CREATE DIRECTORY PLSHPROF_DIR as '/private/plshprof/results';
```

To run the SQL script in Example 8–1, you must have `READ` and `WRITE` privileges on both `PLSHPROF_DIR` and the directory to which it is mapped. if you are connected to the database `AS SYSDBA`, this `GRANT` statement grants `READ` and `WRITE` privileges on `PLSHPROF_DIR` to `HR`:

```
GRANT READ, WRITE ON DIRECTORY PLSHPROF_DIR TO HR;
```

For more information about using directory objects, see *Oracle Database SecureFiles and Large Objects Developer's Guide*.

Understanding Raw Profiler Output

Raw profiler output is intended to be processed by the analyzer component of the PL/SQL hierarchical profiler. However, even without such processing, it provides a

simple function-level trace of the program. This topic explains how to understand raw profiler output.

> **Note:** The raw profiler format shown in this chapter is intended only to illustrate conceptual features of raw profiler output. Format specifics are subject to change at each Oracle Database release.

The SQL script in Example 8–1 wrote this raw profiler output to the file `test.trc`:

```
P#V PLSHPROF Internal Version 1.0
P#! PL/SQL Timer Started
P#C PLSQL."".""."__plsql_vm"
P#X 2
P#C PLSQL."".""."__anonymous_block"
P#X 50
P#C PLSQL."HR"."TEST"::7."TEST"#980980e97e42f8ec #1
P#X 3
P#C PLSQL."HR"."TEST"::7."TEST.FOO"#980980e97e42f8ec #4
P#X 35
P#C SQL."HR"."TEST"::7."__static_sql_exec_line6" #6
P#X 206
P#R
P#X 26
P#R
P#X 2
P#C PLSQL."HR"."TEST"::7."TEST.FOO"#980980e97e42f8ec #4
P#X 4
P#C SQL."HR"."TEST"::7."__static_sql_exec_line6" #6
P#X 80
P#R
P#X 3
P#R
P#X 0
P#C PLSQL."HR"."TEST"::7."TEST.FOO"#980980e97e42f8ec #4
P#X 3
P#C SQL."HR"."TEST"::7."__static_sql_exec_line6" #6
P#X 69
P#R
P#X 3
P#R
P#X 1
P#R
P#X 1
P#R
P#X 3
P#R
P#C PLSQL."".""."__plsql_vm"
P#X 3
P#C PLSQL."".""."__anonymous_block"
P#X 44
P#C PLSQL."SYS"."DBMS_HPROF"::11."STOP_PROFILING"#980980e97e42f8ec #53
P#R
P#R
P#R
P#! PL/SQL Timer Stopped
```

Table 8–1 Raw Profiler Output File Indicators

Indicator	Meaning
`P#V`	`PLSHPROF` banner with version number
`P#C`	Call to a subprogram (call event)
`P#R`	Return from a subprogram (return event)
`P#X`	Elapsed time between preceding and following events
`P#!`	Comment

Call events (`P#C`) and return events (`P#R`) are always properly nested (like matched parentheses). If an unhandled exception causes a called subprogram to exit, the profiler still reports a matching return event.

Each call event (`P#C`) entry in the raw profiler output includes this information:

- **Namespace** to which the called subprogram belongs

  See "Namespaces of Tracked Subprograms" on page 8-6.

- **Name of the PL/SQL module** in which the called subprogram is defined
- **Type of the PL/SQL module** in which the called subprogram is defined
- **Name of the called subprogram**

  This name might be one of the "Special Function Names" on page 8-6.

- Hexadecimal value that represents an MD5 hash of the **signature** of the called subprogram

  The `DBMS_HPROF.analyze` PL/SQL API (described in "Analyzing Profile Data" on page 8-6) stores the hash value in a hierarchical profiler table, which allows both you and `DBMS_HPROF.analyze` to distinguish between **overloaded subprograms** (subprograms with the same name).

- **Line number** at which the called subprogram is defined in the PL/SQL module

  PL/SQL hierarchical profiler does not measure time spent at individual lines within modules, but you can use line numbers to identify the source locations of subprograms in the module (as IDE/Tools do) and to distinguish between overloaded subprograms.

For example, consider this entry in the preceding example of raw profiler output:

```
P#C PLSQL."HR"."TEST"::7."TEST.FOO"#980980e97e42f8ec #4
```

The components of the preceding entry have these meanings:

Component	Meaning
`PLSQL`	`PLSQL` is the **namespace** to which the called subprogram belongs.
`"HR"."TEST"`	`HR.TEST` is the **name of the PL/SQL module** in which the called subprogram is defined.
7	7 is the internal enumerator for the **module type** of `HR.TEST`. Examples of module types are procedure, package, and package body.
`"TEST.FOO"`	`TEST.FOO` is the **name of the called subprogram**.
`#980980e97e42f8ec`	`#980980e97e42f8ec` is a hexadecimal value that represents an MD5 hash of the **signature** of `TEST.FOO`.

Component	Meaning
#4	4 is the **line number** in the PL/SQL module HR.TEST at which TEST.FOO is defined.

Note: When a subprogram is inlined, it is not reported in the profiler output. For information about subprogram inlining, see *Oracle Database PL/SQL Language Reference*.

When a call to a DETERMINISTIC function is "optimized away," it is not reported in the profiler output. For information about DETERMINISTIC functions, see *Oracle Database PL/SQL Language Reference*.

Namespaces of Tracked Subprograms

The subprograms that the PL/SQL hierarchical profiler tracks are classified into the namespaces PLSQL and SQL, as follows:

- Namespace PLSQL includes:
  - PL/SQL subprogram calls
  - PL/SQL triggers
  - PL/SQL anonymous blocks
  - Remote subprogram calls
  - Package initialization blocks
- Namespace SQL includes SQL statements executed from PL/SQL, such as queries, data manipulation language (DML) statements, data definition language (DDL) statements, and native dynamic SQL statements

Special Function Names

PL/SQL hierarchical profiler tracks certain operations as if they were functions with the names and namespaces shown in Table 8–2.

Table 8–2 Function Names of Operations that the PL/SQL Hierarchical Profiler Tracks

Tracked Operation	Function Name	Namespace
Call to PL/SQL Virtual Machine	__plsql_vm	PL/SQL
PL/SQL anonymous block	__anonymous_block	PL/SQL
Package initialization block	__pkg_init	PL/SQL
Static SQL statement at line *line#*	__static_sql_exec_line*line#*	SQL
Dynamic SQL statement at line *line#*	__dyn_sql_exec_line*line#*	SQL
SQL FETCH statement at line *line#*	__sql_fetch_line*line#*	SQL

Analyzing Profile Data

The analyzer component of the PL/SQL hierarchical profiler, DBMS_HPROF.analyze, processes the raw profiler output and stores the results in the **hierarchical database tables** described in Table 8–3.

Table 8–3 PL/SQL Hierarchical Profiler Database Tables

Table	Description
DBMSHP_RUNS	Top-level information for this run of DBMS_HPROF.analyze. For column descriptions, see Table 8–4 on page 8-8.
DBMSHP_FUNCTION_INFO	Information for each subprogram profiled in this run of DBMS_HPROF.analyze. For column descriptions, see Table 8–5 on page 8-9.
DBMSHP_PARENT_CHILD_INFO	Parent-child information for each subprogram profiled in this run of DBMS_HPROF.analyze. For column descriptions, see Table 8–6 on page 8-8.

Topics:

- Creating Hierarchical Profiler Tables
- Understanding Hierarchical Profiler Tables

Note: To generate simple HTML reports directly from raw profiler output, without using the Analyzer, you can use the plshprof command-line utility. For details, see "plshprof Utility" on page 8-13.

Creating Hierarchical Profiler Tables

To create the hierarchical profiler tables in Table 8–3 and the other data structures required for persistently storing profile data, follow these steps:

1. Run the script dbmshptab.sql (located in the directory rdbms/admin).

   This script creates both the hierarchical profiler tables in Table 8–3 and the other data structures required for persistently storing profile data.

   Note: Running the script dbmshptab.sql drops any previously created hierarchical profiler tables.

2. Ensure that you have these privileges:
   - EXECUTE privilege on the DBMS_HPROF package
   - READ privilege on the directory that DBMS_HPROF.analyze specifies
3. Use the PL/SQL API DBMS_HPROF.analyze to analyze a single raw profiler output file and store the results in hierarchical profiler tables.

   (For an example of a raw profiler output file, see test.trc in "Understanding Raw Profiler Output" on page 8-3.)

   For more information about DBMS_HPROF.analyze, see *Oracle Database PL/SQL Packages and Types Reference*.
4. Use the hierarchical profiler tables to generate custom reports.

The anonymous block in Example 8–2:

- Invokes the function DBMS_HPROF.analyze function, which:
  - Analyzes the profile data in the raw profiler output file test.trc (from "Understanding Raw Profiler Output" on page 8-3), which is in the directory

that is mapped to the directory object PLSHPROF_DIR, and stores the data in the hierarchical profiler tables in Table 8–3 on page 8-7.

  - Returns a unique identifier that you can use to query the hierarchical profiler tables in Table 8–3 on page 8-7. (By querying these hierarchical profiler tables, you can produce customized reports.)

- Stores the unique identifier in the variable runid, which it prints.

Example 8–2 Invoking DBMS_HPROF.analyze

```
DECLARE
  runid NUMBER;
BEGIN
  runid := DBMS_HPROF.analyze(LOCATION=>'PLSHPROF_DIR',
                              FILENAME=>'test.trc');
  DBMS_OUTPUT.PUT_LINE('runid = ' || runid)
END;
/
```

Understanding Hierarchical Profiler Tables

These topics explain how to use the hierarchical profiler tables in Table 8–3:

- Hierarchical Profiler Database Table Columns
- Distinguishing Between Overloaded Subprograms
- Hierarchical Profiler Tables for Sample PL/SQL Procedure
- Examples of Calls to DBMS_HPROF.analyze with Options

Hierarchical Profiler Database Table Columns

Table 8–4 describes the columns of the hierarchical profiler table DBMSHP_RUNS, which contains one row of top-level information for each run of DBMS_HPROF.analyze.

The primary key for the hierarchical profiler table DBMSHP_RUNS is RUNID.

Table 8–4 DBMSHP_RUNS Table Columns

Column Name	Column Data Type	Column Contents
RUNID	NUMBER PRIMARY KEY	Unique identifier for this run of DBMS_HPROF.analyze, generated from DBMSHP_RUNNUMBER sequence.
RUN_TIMESTAMP	TIMESTAMP	Time stamp for this run of DBMS_HPROF.analyze.
RUN_COMMENT	VARCHAR2(2047)	Comment that you provided for this run of DBMS_HPROF.analyze.
TOTAL_ELAPSED_TIME	INTEGER	Total elapsed time for this run of DBMS_HPROF.analyze.

Table 8–5 describes the columns of the hierarchical profiler table DBMSHP_FUNCTION_INFO, which contains one row of information for each subprogram profiled in this run of DBMS_HPROF.analyze. If a subprogram is overloaded, Table 8–5 has a row for each variation of that subprogram. Each variation has its own LINE# and HASH (see "Distinguishing Between Overloaded Subprograms" on page 8-10).

The primary key for the hierarchical profiler table DBMSHP_FUNCTION_INFO is RUNID, SYMBOLID.

Table 8–5 DBMSHP_FUNCTION_INFO Table Columns

Column Name	Column Data Type	Column Contents
RUNID	NUMBER	References RUNID column of DBMSHP_RUNS table. For a description of that column, see Table 8–4.
SYMBOLID	NUMBER	Symbol identifier for subprogram (unique for this run of DBMS_HPROF.analyze).
OWNER	VARCHAR2(32)	Owner of module in which each subprogram is defined (for example, SYS or HR).
MODULE	VARCHAR2(2047)	Module in which subprogram is defined (for example, DBMS_LOB, UTL_HTTP, or MY_PACKAGE).
TYPE	VARCHAR2(32)	Type of module in which subprogram is defined (for example, PACKAGE, PACKAGE_BODY, or PROCEDURE).
FUNCTION	VARCHAR2(4000)	Name of subprogram (not necessarily a function) (for example, INSERT_ORDER, PROCESS_ITEMS, or TEST). This name might be one of the "Special Function Names" on page 8-6. For subprogram B defined within subprogram A, this name is A.B. A recursive call to function X is denoted X@*n*, where *n* is the recursion depth. For example, X@1 is the first recursive call to X.
LINE#	NUMBER	Line number in OWNER.MODULE at which FUNCTION is defined.
HASH	RAW(32)	Hash code for signature of subprogram (unique for this run of DBMS_HPROF.analyze).
NAMESPACE	VARCHAR2(32)	Namespace of subprogram. For details, see "Namespaces of Tracked Subprograms" on page 8-6.
SUBTREE_ELAPSED_TIME	INTEGER	Elapsed time, in microseconds, for subprogram, including time spent in descendant subprograms.
FUNCTION_ELAPSED_TIME	INTEGER	Elapsed time, in microseconds, for subprogram, excluding time spent in descendant subprograms.
CALLS	INTEGER	Number of calls to subprogram.

Table 8–6 describes the columns of the hierarchical profiler table DBMSHP_PARENT_CHILD_INFO, which contains one row of parent-child information for each unique parent-child subprogram combination profiled in this run of DBMS_HPROF.analyze.

Table 8–6 DBMSHP_PARENT_CHILD_INFO Table Columns

Column Name	Column Data Type	Column Contents
RUNID	NUMBER	References RUNID column of DBMSHP_FUNCTION_INFO table. For a description of that column, see Table 8–5.
PARENTSYMID	NUMBER	Parent symbol ID. RUNID, PARENTSYMID references DBMSHP_FUNCTION_INFO(RUNID, SYMBOLID).
CHILDSYMID	VARCHAR2(32)	Child symbol ID. RUNID, CHILDSYMID references DBMSHP_FUNCTION_INFO(RUNID, SYMBOLID).
SUBTREE_ELAPSED_TIME	INTEGER	Elapsed time, in microseconds, for RUNID, CHILDSYMID when called from RUNID, PARENTSYMID, including time spent in descendant subprograms.
FUNCTION_ELAPSED_TIME	INTEGER	Elapsed time, in microseconds, for RUNID, CHILDSYMID when called from RUNID, PARENTSYMID, excluding time spent in descendant subprograms.
CALLS	INTEGER	Number of calls to RUNID, CHILDSYMID from RUNID, PARENTSYMID.

Distinguishing Between Overloaded Subprograms

Overloaded subprograms are different subprograms with the same name (see *Oracle Database PL/SQL Language Reference*).

Suppose that a program declares three subprograms named compute—the first at line 50, the second at line 76, and the third at line 100. In the DBMSHP_FUNCTION_INFO table, each of these subprograms has compute in the FUNCTION column. To distinguish between the three subprograms, use either the LINE# column (which has 50 for the first subprogram, 76 for the second, and 100 for the third) or the HASH column (which has a unique value for each subprogram).

In the profile data for two different runs, the LINE# and HASH values for a function might differ. The LINE# value of a function changes if you insert or delete lines before the function definition. The HASH value changes only if the signature of the function changes; for example, if you change the parameter list.

Hierarchical Profiler Tables for Sample PL/SQL Procedure

The hierarchical profiler tables for the PL/SQL procedure test in "Collecting Profile Data" on page 8-2 are shown in Example 8–3 through Example 8–5.

Example 8–3 DBMSHP_RUNS Table for Sample PL/SQL Procedure

```
RUNID  RUN_TIMESTAMP                TOTAL_ELAPSED_TIME  RUN_COMMENT
1      10-APR-06 12.01.56.766743 PM  2637                 First run of TEST
```

Example 8–4 DBMSHP_FUNCTION_INFO Table for Sample PL/SQL Procedure

RUNID	SYMBOLID	OWNER	MODULE	TYPE	NAMESPACE	FUNCTION
1	1				PLSQL	__anonymous_block
1	2				PLSQL	__plsql_vm
1	3	HR	TEST	PROCEDURE	PLSQL	TEST
1	4	HR	TEST	PROCEDURE	PLSQL	TEST.FOO
1	5	SYS	DBMS_HPROF	PACKAGE_BODY	PLSQL	STOP_PROFILING
1	6	HR	TEST	PROCEDURE	SQL	__static_sql_exec_line5

LINE#	CALLS	HASH	SUBTREE_ELAPSED_TIME	FUNCTION_ELAPSED_TIME
0	2	980980E97E42F8EC	2554	342
0	2	980980E97E42F8EC	2637	83
1	1	980980E97E42F8EC	2212	28
3	3	980980E97E42F8EC	2184	126
57	1	980980E97E42F8EC	0	0
5	3	980980E97E42F8EC	1998	1998

Example 8–5 DBMSHP_PARENT_CHILD_INFO Table for Sample PL/SQL Procedure

RUNID	PARENTSYMID	CHILDSYMID	SUBTREE_ELAPSED_TIME	FUNCTION_ELAPSED_TIME	CALLS
1	2	1	2554	342	2
1	1	3	2212	28	1
1	3	4	2184	126	3
1	1	5	0	0	1
1	4	6	1998	1998	3

Consider the third row of the table DBMSHP_PARENT_CHILD_INFO (Example 8–5). The RUNID column shows that this row corresponds to the first run. The columns PARENTSYMID and CHILDSYMID show that the symbol IDs of the parent (caller) and child (called subprogram) are 3 and 4, respectively. The table DBMSHP_FUNCTION_INFO (Example 8–4) shows that for the first run, the symbol IDs 3 and 4 correspond to procedures TEST and TEST.FOO, respectively. Therefore, the information in this row is about calls from the procedure TEST to the procedure FOO (defined within TEST) in the module HR.TEST. This row shows that, when called from the procedure TEST, the function time for the procedure FOO is 126 microseconds, and the time spent in the FOO subtree (including descendants) is 2184 microseconds.

Examples of Calls to DBMS_HPROF.analyze with Options

For an example of a call to DBMS_HPROF.analyze without options, see Example 8–2 on page 8-8.

Example 8–6 creates a package, creates a procedure that invokes subprograms in the package, profiles the procedure, and uses DBMS_HRPROF.analyze to analyze the raw profiler output file. The raw profiler output file is in the directory corresponding to the PLSHPROF_DIR directory object.

Example 8–6 Invoking DBMS_HPROF.analyze with Options

```
-- Create package

CREATE OR REPLACE PACKAGE pkg IS
  PROCEDURE myproc (n IN out NUMBER);
  FUNCTION myfunc (v VARCHAR2) RETURN VARCHAR2;
  FUNCTION myfunc (n PLS_INTEGER) RETURN PLS_INTEGER;
END pkg;
/
CREATE OR REPLACE PACKAGE BODY pkg IS
  PROCEDURE myproc (n IN OUT NUMBER) IS
```

```
  BEGIN
    n := n + 5;
  END;

  FUNCTION myfunc (v VARCHAR2) RETURN VARCHAR2 IS
    n NUMBER;
  BEGIN
    n := LENGTH(v);
    myproc(n);
    IF n > 20 THEN
      RETURN SUBSTR(v, 1, 20);
    ELSE
      RETURN v || '...';
    END IF;
  END;

  FUNCTION myfunc (n PLS_INTEGER) RETURN PLS_INTEGER IS
    i PLS_INTEGER;
    PROCEDURE myproc (n IN out PLS_INTEGER) IS
    BEGIN
      n := n + 1;
    END;
  BEGIN
    i := n;
    myproc(i);
    RETURN i;
  END;
END pkg;
/

-- Create procedure that invokes package subprograms

CREATE OR REPLACE PROCEDURE test2 IS
  x NUMBER := 5;
  y VARCHAR2(32767);
BEGIN
  pkg.myproc(x);
  y := pkg.myfunc('hello');
END;

-- Profile test2

BEGIN
  DBMS_HPROF.START_PROFILING('PLSHPROF_DIR', 'test2.trc');
END;
/
BEGIN
  test2;
END;
/
BEGIN
  DBMS_HPROF.STOP_PROFILING;
END;
/
-- If not done, create hierarchical profiler tables
-- (see "Creating Hierarchical Profiler Tables" on page 8-7.)

-- Call DBMS_HPROF.analyze with options

DECLARE
```

```
  runid NUMBER;
BEGIN
  -- Analyze only subtrees rooted at trace entry "HR"."PKG"."MYPROC"

  runid := DBMS_HPROF.analyze('PLSHPROF_DIR', 'test2.trc',
                              trace => '"HR"."PKG"."MYPROC"');

  -- Analyze up to 20 calls to subtrees rooted at trace entry
  -- "HR"."PKG"."MYFUNC".  Because "HR"."PKG"."MYFUNC" is overloaded,
  -- both overloads are considered for analysis.

  runid := DBMS_HPROF.analyze('PLSHPROF_DIR', 'test2.trc',
                              collect => 20,
                              trace => '"HR"."PKG"."MYFUNC"');

  -- Analyze second call to PL/SQL virtual machine

  runid := DBMS_HPROF.analyze('PLSHPROF_DIR', 'test2.trc',
                              skip => 1, collect => 1,
                              trace => """"."""__plsql_vm"');
END;
/
```

plshprof Utility

The `plshprof` command-line utility (located in the directory `$ORACLE_HOME/bin/`) generates simple HTML reports from either one or two raw profiler output files. (For an example of a raw profiler output file, see `test.trc` in "Collecting Profile Data" on page 8-2.)

You can browse the generated HTML reports in any browser. The browser's navigational capabilities, combined with well chosen links, provide a powerful way to analyze performance of large applications, improve application performance, and lower development costs.

Topics:

- plshprof Options
- HTML Report from a Single Raw Profiler Output File
- HTML Difference Report from Two Raw Profiler Output Files

plshprof Options

The command to run the `plshprof` utility is:

```
plshprof [option...] profiler_output_filename_1 profiler_output_filename_2
```

Each *`option`* is one of these:

Option	Description	Default
`-skip` *`count`*	Skips first *`count`* calls. Use only with `-trace` *`symbol`*.	0
`-collect` *`count`*	Collects information for *`count`* calls. Use only with `-trace` *`symbol`*.	1
`-output` *`filename`*	Specifies name of output file	*`symbol`*`.html` or *`tracefile1`*`.html`

Option	Description	Default
`-summary`	Prints only elapsed time	None
`-trace` *`symbol`*	Specifies function name of tree root	Not applicable

Suppose that your raw profiler output file, `test.trc`, is in the current directory. You want to analyze and generate HTML reports, and you want the root file of the HTML report to be named `report.html`. Use this command (`%` is the prompt):

```
% plshprof -output report test.trc
```

HTML Report from a Single Raw Profiler Output File

To generate a PL/SQL hierarchical profiler HTML report from a single raw profiler output file, use these commands:

```
% cd target_directory
% plshprof -output html_root_filename profiler_output_filename
```

`target_directory` is the directory in which you want the HTML files to be created.

`html_root_filename` is the name of the root HTML file to be created.

`profiler_output_filename` is the name of a raw profiler output file.

The preceding `plshprof` command generates a set of HTML files. Start browsing them from *`html_root_filename`*`.html`.

Topics:

- First Page of Report
- Function-Level Reports
- Module-Level Reports
- Namespace-Level Reports
- Parents and Children Report for a Function

First Page of Report

The first page of an HTML report from a single raw profiler output file includes summary information and hyperlinks to other pages of the report.

Sample First Page

PL/SQL Elapsed Time (microsecs) Analysis

2831 microsecs (elapsed time) & 12 function calls

The PL/SQL Hierarchical Profiler produces a collection of reports that present information derived from the profiler output log in a variety of formats. These reports have been found to be the most generally useful as starting points for browsing:

- Function Elapsed Time (microsecs) Data sorted by Total Subtree Elapsed Time (microsecs)
- Function Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)

In addition, the following reports are also available:

- Function Elapsed Time (microsecs) Data sorted by Function Name

- Function Elapsed Time (microsecs) Data sorted by Total Descendants Elapsed Time (microsecs)
- Function Elapsed Time (microsecs) Data sorted by Total Function Call Count
- Function Elapsed Time (microsecs) Data sorted by Mean Subtree Elapsed Time (microsecs)
- Function Elapsed Time (microsecs) Data sorted by Mean Function Elapsed Time (microsecs)
- Function Elapsed Time (microsecs) Data sorted by Mean Descendants Elapsed Time (microsecs)
- Module Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)
- Module Elapsed Time (microsecs) Data sorted by Module Name
- Module Elapsed Time (microsecs) Data sorted by Total Function Call Count
- Namespace Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)
- Namespace Elapsed Time (microsecs) Data sorted by Namespace
- Namespace Elapsed Time (microsecs) Data sorted by Total Function Call Count
- Parents and Children Elapsed Time (microsecs) Data

Function-Level Reports

The function-level reports provide a flat view of the profile information. Each function-level report includes this information for each function:

- Function time (time spent in the function itself, also called "self time")
- Descendants time (time spent in the descendants of the function)
- Subtree time (time spent in the subtree of the function—function time plus descendants time)
- Number of calls to the function
- Function name

  The function name is hyperlinked to the Parents and Children Report for the function.

Each function-level report is sorted on a particular attribute; for example, function time or subtree time.

This sample report is sorted in descending order of the total subtree elapsed time for the function, which is why information in the Subtree and Ind% columns is in **bold type**:

Sample Report

Function Elapsed Time (microsecs) Data sorted by Total Subtree Elapsed Time (microsecs)

2831 microsecs (elapsed time) & 12 function calls

Subtree	Ind%	Function	Descendant	Ind%	Calls	Ind%	Function Name
2831	**100%**	93	2738	96.7%	2	16.7%	__plsq_vm

Subtree	Ind%	Function	Descendant	Ind%	Calls	Ind%	Function Name
2738	**96.7%**	310	2428	85.8%	2	16.7%	__anonymous_block
2428	**85.8%**	15	2413	85.2%	1	8.3%	HR.TEST.TEST (Line 1)
2413	**85.2%**	435	1978	69.9%	3	25.0%	HR.TEST.TEST.FOO (Line 3)
1978	**69.9%**	1978	0	0.0%	3	25.0%	HR.TEST.__static_sql_exec_ line5 (Line 5)
0	**0.0%**	0	0	0.0%	1	8.3%	SYS.DBMS_HPROF.STOP_ PROFILING (Line 53)

Module-Level Reports

Each module-level report includes this information for each module (for example, package or type):

- Module time (time spent in the module—sum of the function times of all functions in the module)
- Number of calls to functions in the module

Each module-level report is sorted on a particular attribute; for example, module time or module name.

This sample report is sorted by module time, which is why information in the Module, Ind%, and Cum% columns is in **bold type**:

Sample Report

Module Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)

166878 microsecs (elapsed time) & 1099 function calls

Module	Ind%	Cum%	Calls	Ind%	Module Name
84932	**50.9%**	**50.9%**	6	0.5%	HR.P
67749	**40.6%**	**91.5%**	216	19.7%	SYS.DBMS_LOB
13340	**8.0%**	**99.5%**	660	60.1%	SYS.UTL_FILE
839	**0.5%**	**100%**	214	19.5%	SYS.UTL_RAW
18	**0.0%**	**100%**	2	0.2%	HR.UTILS
0	**0.0%**	**100%**	1	0.1%	SYS.DBMS_HPROF

Namespace-Level Reports

Each namespace-level report includes this information for each namespace:

- Namespace time (time spent in the namespace—sum of the function times of all functions in the namespace)
- Number of calls to functions in the namespace

Each namespace-level report is sorted on a particular attribute; for example, namespace time or number of calls to functions in the namespace.

This sample report is sorted by function time, which is why information in the Function, Ind%, and Cum% columns is in **bold type**:

Sample Report

Namespace Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)

166878 microsecs (elapsed time) & 1099 function calls

Function	Ind%	Cum%	Calls	Ind%	Namespace
93659	**56.1%**	**56.1%**	1095	99.6%	PLSQL
73219	**43.9%**	**100%**	4	0.4%	SQL

Parents and Children Report for a Function

For each function tracked by the profiler, the Parents and Children Report provides information about parents (functions that call it) and children (functions that it calls). For each parent, the report gives the function's execution profile (subtree time, function time, descendants time, and number of calls). For each child, the report gives the execution profile for the child when called from this function (but not when called from other functions).

The execution profile for a function includes the same information for that function as a function-level report includes for each function (for details, see "Function-Level Reports" on page 8-15).

This Sample Report is a fragment of a Parents and Children Report that corresponds to a function named `HR.P.UPLOAD`. The first row has this summary information:

- There are two calls to the function `HR.P.UPLOAD`.
- The total subtree time for the function is 166,860 microseconds—11,713 microseconds (7.0%) in the function itself and 155,147 microseconds (93.0%) in its descendants.

After the row "Parents" are the execution profile rows for the two parents of `HR.P.UPLOAD`, which are `HR.UTILS.COPY_IMAGE` and `HR.UTILS.COPY_FILE`.

The first parent execution profile row, for `HR.UTILS.COPY_IMAGE`, shows:

- `HR.UTILS.COPY_IMAGE` calls `HR.P.UPLOAD` once, which is 50% of the number of calls to `HR.P.UPLOAD`.
- The subtree time for `HR.P.UPLOAD` when called from `HR.UTILS.COPY_IMAGE` is 106,325 microseconds, which is 63.7% of the total subtree time for `HR.P.UPLOAD`.
- The function time for `HR.P.UPLOAD` when called from `HR.UTILS.COPY_IMAGE` is 6,434 microseconds, which is 54.9% of the total function time for `HR.P.UPLOAD`.

After the row "Children" are the execution profile rows for the children of `HR.P.UPLOAD` when called from `HR.P.UPLOAD`.

The third child execution profile row, for `SYS.UTL_FILE.GET_RAW`, shows:

- `HR.P.UPLOAD` calls `SYS.UTL_FILE.GET_RAW` 216 times.
- The subtree time, function time and descendants time for `SYS.UTL_FILE.GET_RAW` when called from `HR.P.UPLOAD` are 12,487 microseconds, 3,969 microseconds, and 8,518 microseconds, respectively.
- Of the total descendants time for `HR.P.UPLOAD` (155,147 microseconds), the child `SYS.UTL_FILE.GET_RAW` is responsible for 12,487 microsecs (8.0%).

Sample Report

HR.P.UPLOAD (Line 3)

Subtree	Ind%	Function	Ind%	Descendant	Ind%	Calls	Ind%	Function Name
166860	**100%**	11713	7.0%	155147	93.0%	2	0.2%	HR.P.UPLOAD (Line 3)
Parents:								
106325	**63.7%**	6434	54.9%	99891	64.4%	1	50.0%	HR.UTILS.COPY_IMAGE (Line 3)
60535	**36.3%**	5279	45.1%	55256	35.6%	1	50.0%	HR.UTILS.COPY_FILE (Line 8))
Children:								
71818	**46.3%**	71818	100%	0	N/A	2	100%	HR.P.__static_sql_exec_line38 (Line 38)
67649	**43.6%**	67649	100%	0	N/A	214	100%	SYS.DBMS_LOB.WRITEAPPEND (Line 926)
12487	**8.0%**	3969	100%	8518	100%	216	100%	SYS.UTL_FILE.GET_RAW (Line 1089)
1401	**0.9%**	1401	100%	0	N/A	2	100%	HR.P.__static_sql_exec_line39 (Line 39)
839	**0.5%**	839	100%	0	N/A	214	100%	SYS.UTL_FILE.GET_RAW (Line 246)
740	**0.5%**	73	100%	667	100%	2	100%	SYS.UTL_FILE.FOPEN (Line 422)
113	**0.1%**	11	100%	102	100%	2	100%	SYS.UTL_FILE.FCLOSE (Line 585)
100	**0.1%**	100	100%	0	N/A	2	100%	SYS.DBMS_LOB.CREATETEMPORARY (Line 536)

HTML Difference Report from Two Raw Profiler Output Files

To generate a PL/SQL hierarchical profiler HTML difference report from two raw profiler output files, use these commands:

```
% cd target_directory
% plshprof -output html_root_filename profiler_output_filename_1 profiler_output_filename_2
```

target_directory is the directory in which you want the HTML files to be created.

html_root_filename is the name of the root HTML file to be created.

profiler_output_filename_1 and *profiler_output_filename_2* are the names of raw profiler output files.

The preceding `plshprof` command generates a set of HTML files. Start browsing them from *html_root_filename*.html.

Topics:

- Difference Report Conventions

- First Page of Difference Report
- Function-Level Difference Reports
- Module-Level Difference Reports
- Namespace-Level Difference Reports
- Parents and Children Difference Report for a Function

Difference Report Conventions

Difference reports use these conventions:

- In a report title, **Delta** means **difference**, or **change**.
- A **positive value** indicates that the number increased (**regressed**) from the first run to the second run.
- A **negative value** for a difference indicates that the number decreased (**improved**) from the first run to the second run.
- The symbol **#** after a function name means that the function was called in only one run.

First Page of Difference Report

The first page of an HTML difference report from two raw profiler output files includes summary information and hyperlinks to other pages of the report.

Sample First Page

PL/SQL Elapsed Time (microsecs) Analysis – Summary Page

This analysis finds a net **regression** of **2709589** microsecs (elapsed time) or **80%** (**3393719** versus **6103308**). Here is a summary of the 7 most important individual function regressions and improvements:

Regressions: 3399382 microsecs (elapsed time)

Function	Rel%	Ind%	Calls	Rel%	Function Name
2075627	**+941%**	**61.1%**	0		HR.P.G (Line 35)
1101384	**+54.6%**	**32.4%**	5	+55.6%	HR.P.H (Line 18)
222371		**6.5%**	1		HR.P.J (Line 10)

Improvements: 689793 microsecs (elapsed time)

Function	Rel%	Ind%	Calls	Rel%	Function Name
-467051	**-50.0%**	**67.7%**	-2	-50.0%	HR.P.F (Line 25)
-222737		**32.3%**	-1		HR.P.I (Line 2)#
-5	**-21.7%**	**0.0%**	0		HR.P.TEST (Line 46)

The PL/SQL Timing Analyzer produces a collection of reports that present information derived from the profiler's output logs in a variety of formats. The following reports have been found to be the most generally useful as starting points for browsing:

- Function Elapsed Time (microsecs) Data for Performance Regressions

- Function Elapsed Time (microsecs) Data for Performance Improvements

In addition, the following reports are also available:

- Function Elapsed Time (microsecs) Data sorted by Function Name
- Function Elapsed Time (microsecs) Data sorted by Total Subtree Elapsed Time (microsecs) Delta
- Function Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs) Delta
- Function Elapsed Time (microsecs) Data sorted by Total Descendants Elapsed Time (microsecs) Delta
- Function Elapsed Time (microsecs) Data sorted by Total Function Call Count Delta
- Module Elapsed Time (microsecs) Data sorted by Module Name
- Module Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs) Delta
- Module Elapsed Time (microsecs) Data sorted by Total Function Call Count Delta
- Namespace Elapsed Time (microsecs) Data sorted by Namespace
- Namespace Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs)
- Namespace Elapsed Time (microsecs) Data sorted by Total Function Call Count
- File Elapsed Time (microsecs) Data Comparison with Parents and Children

Function-Level Difference Reports

Each function-level difference report includes, for each function, the change in these values from the first run to the second run:

- Function time (time spent in the function itself, also called "self time")
- Descendants time (time spent in the descendants of the function)
- Subtree time (time spent in the subtree of the function—function time plus descendants time)
- Number of calls to the function
- Mean function time

  The mean function time is the function time divided by number of calls to the function.
- Function name

  The function name is hyperlinked to the Parents and Children Difference Report for the function.

The report in Sample Report 1 shows the difference information for all functions that performed better in the first run than they did in the second run. Note that:

- For `HR.P.G`, the function time increased by 2,075,627 microseconds (941%), which accounts for 61.1% of all regressions.
- For `HR.P.H`, the function time and number of calls increased by 1,101,384 microseconds (54.6%) and 5 (55.6%), respectively, but the mean function time improved by 1,346 microseconds (-0.6%).
- `HR.P.J` was called in only one run.

Sample Report 1

Function Elapsed Time (microsecs) Data for Performance Regressions

Subtree	Function	Rel%	Ind%	Cum%	Descendant	Calls	Rel%	Mean Function	Rel%	Function Name
4075787	**2075627**	**+941%**	**61.1%**	61.1%	2000160	0		2075627	+941%	HR.P.G (Line 35)
1101384	**1101384**	**+54.6%**	**32.4%**	93.5%	0	5	+55.6%	-1346	-0.6%	HR.P.H (Line 18)
222371	**222371**		6.5%	100%	0	1				HR.P.J (Line 10)#

The report in Sample Report 2 shows the difference information for all functions that performed better in the second run than they did in the first run.

Sample Report 2

Function Elapsed Time (microsecs) Data for Performance Improvements

Subtree	Function	Rel%	Ind%	Cum%	Descendant	Calls	Rel%	Mean Function	Rel%	Function Name
-1365827	**-467051**	**-50.0%**	**67.7%**	67.7%	-898776	-2	-50.0%	-32	0.0%	HR.P.F (Line 25)
-222737	**-222737**		**32.3%**	100%	0	-1				HR.P.I (Line 2)
2709589	**-5**	**-21.7%**	**0.0%**	100%	2709594	0		-5	-20.8	HR.P.TEST (Line 46)#

The report in Sample Report 3 summarizes the difference information for all functions.

Sample Report 3

Function Elapsed Time (microsecs) Data sorted by Total Function Call Count Delta

Subtree	Function	Rel%	Ind%	Descendant	Calls	Rel%	Mean Function	Rel%	Function Name
1101384	1101384	+54.6%	32.4%	0	**5**	**+55.6%**	-1346	-0.6%	HR.P.H (Line 18)
-1365827	-467051	+50.0%	67.7%	-898776	**-2**	**-50.0%**	-32	-0.0%	HR.P.F (Line 25)
-222377	-222377		32.3%	0	**-1**				HR.P.I (Line 2)#
222371	222371		6.5%	0	**1**				HR.P.J(Line 10)#
4075787	2075627	+941%	61.1%	2000160	**0**		2075627	+941%	HR.P.G (Line 35)
2709589	-5	-21.7%	0.0%	2709594	**0**		-5	-20.8%	HR.P.TEST (Line 46)
0	0			0	**0**				SYS.DBMS_HPROF.STOP_ PROFILING (Line 53)

Module-Level Difference Reports

Each module-level report includes, for each module, the change in these values from the first run to the second run:

- Module time (time spent in the module—sum of the function times of all functions in the module)
- Number of calls to functions in the module

Sample Report

Module Elapsed Time (microsecs) Data sorted by Total Function Elapsed Time (microsecs) Delta

Module	Calls	Module Name
2709589	3	HR.P

Module	Calls	Module Name
0	0	SYS.DBMS_HPROF

Namespace-Level Difference Reports

Each namespace-level report includes, for each namespace, the change in these values from the first run to the second run:

- Namespace time (time spent in the namespace—sum of the function times of all functions in the namespace)
- Number of calls to functions in the namespace

Sample Report

Namespace Elapsed Time (microsecs) Data sorted by Namespace

Function	Calls	Namespace
2709589	3	**PLSQL**

Parents and Children Difference Report for a Function

The Parents and Children Difference Report for a function shows changes in the execution profiles of these from the first run to the second run:

- Parents (functions that call the function)
- Children (functions that the function calls)

  Execution profiles for children include only information from when this function calls them, not for when other functions call them.

The execution profile for a function includes this information:

- Function time (time spent in the function itself, also called "self time")
- Descendants time (time spent in the descendants of the function)
- Subtree time (time spent in the subtree of the function—function time plus descendants time)
- Number of calls to the function
- Function name

The sample report is a fragment of a Parents and Children Difference Report that corresponds to a function named `HR.P.X`.

The first row, a summary of the difference between the first and second runs, shows regression: function time increased by 1,094,099 microseconds (probably because the function was called five more times).

The "Parents" rows show that `HR.P.G` called `HR.P.X` nine more times in the second run than it did in the first run, while `HR.P.F` called it four fewer times.

The "Children" rows show that `HR.P.X` called each child five more times in the second run than it did in the first run.

Sample Report

HR.P.X (Line 11)

Subtree	Function	Descendant	Calls	Function Name
3322196	1094099	2228097	5	**HR.P.X (Line 11)**
Parents:				
6037490	1993169	4044321	9	**HR.P.G (Line 38)**
-2715294	-899070	-1816224	-4	**HR.P.F (Line 28)**
Children:				
1125489	1125489	0	5	**HR.P.J (Line 10)**
1102608	1102608	0	5	**HR.P.I (Line 2)**

The Parents and Children Difference Report for a function is accompanied by a Function Comparison Report, which shows the execution profile of the function for the first and second runs and the difference between them. This example is the Function Comparison Report for the function `HR.P.X`:

Sample Report

Elapsed Time (microsecs) for HR.P.X (Line 11) (20.1% of total regression)

HR.P.X (Line 11)	First Trace	Ind%	Second Trace	Ind%	Diff	Diff%
Function Elapsed Time (microsecs)	1999509	26.9%	3093608	24.9%	1094099	+54.7%
Descendants Elapsed Time (microsecs)	4095943	55.1%	6324040	50.9%	2228097	+54.4%
Subtree Elapsed Time (microsecs)	6095452	81.9%	9417648	75.7%	3322196	+54.5%
Function Calls	9	25.0%	14	28.6%	5	+55.6%
Mean Function Elapsed Time (microsecs)	222167.7		220972.0		-1195.7	-0.5%
Mean Descendants Elapsed Time (microsecs)	455104.8		451717.1		-3387.6	-0.7%
Mean Subtree Elapsed Time (microsecs)	677272.4		672689.1		-4583.3	-0.7%

9

Developing PL/SQL Web Applications

This chapter explains how to develop PL/SQL Web applications, which let you make your database available on the intranet.

Topics:

- Overview of PL/SQL Web Applications
- Implementing PL/SQL Web Applications
- Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application
- Using Embedded PL/SQL Gateway
- Generating HTML Output with PL/SQL
- Passing Parameters to PL/SQL Web Applications
- Performing Network Operations in PL/SQL Subprograms

Overview of PL/SQL Web Applications

Typically, a Web application written in PL/SQL is a set of stored subprograms that interact with Web browsers through HTTP. A set of interlinked, dynamically generated HTML pages forms the user interface of a web application.

The program flow of a PL/SQL Web application is similar to that in a CGI PERL script. Developers often use CGI scripts to produce Web pages dynamically, but such scripts are often not optimal for accessing the database. Delivering Web content with PL/SQL stored subprograms provides the power and flexibility of database processing. For example, you can use data manipulation language (DML) statements, dynamic SQL statements, and cursors. You also eliminate the process overhead of forking a new CGI process to handle each HTTP request.

Figure 9–1 illustrates the generic process for a PL/SQL Web application.

Figure 9–1 PL/SQL Web Application

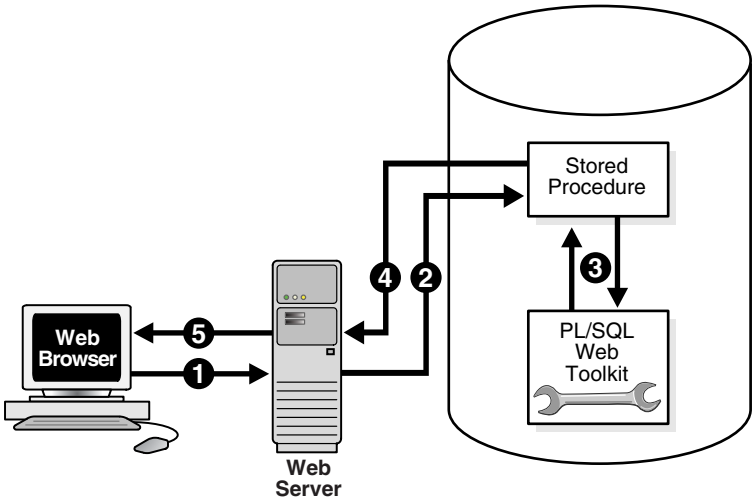

Implementing PL/SQL Web Applications

You can implement a Web browser-based application entirely in PL/SQL with these Oracle Database components:

- PL/SQL Gateway
- PL/SQL Web Toolkit

PL/SQL Gateway

The PL/SQL gateway enables a Web browser to invoke a PL/SQL stored subprogram through an HTTP listener. The gateway is a platform on which PL/SQL users develop and deploy PL/SQL Web applications.

mod_plsql

`mod_plsql` is one implementation of the PL/SQL gateway. The module is a plug-in of Oracle HTTP Server and enables Web browsers to invoke PL/SQL stored subprograms. Oracle HTTP Server is a component of both Oracle Application Server and the database.

The `mod_plsql` plug-in enables you to use PL/SQL stored subprograms to process HTTP requests and generate responses. In this context, an HTTP request is a URL that includes parameter values to be passed to a stored subprogram. PL/SQL gateway translates the URL, invokes the stored subprogram with the parameters, and returns output (typically HTML) to the client.

Some advantages of using `mod_plsql` over the embedded form of the PL/SQL gateway are:

- You can run it in a firewall environment in which the Oracle HTTP Server runs on a firewall-facing host while the database is hosted behind a firewall. You cannot use this configuration with the embedded gateway.
- The embedded gateway does not support `mod_plsql` features such as dynamic HTML caching, system monitoring, and logging in the Common Log Format.

Embedded PL/SQL Gateway

You can use an embedded version of the PL/SQL gateway that runs in the XML DB HTTP Listener in the database. It provides the core features of `mod_plsql` in the database but does not require the Oracle HTTP Server. You configure the embedded PL/SQL gateway with the `DBMS_EPG` package in the PL/SQL Web Toolkit.

Some advantages of using the embedded gateway over `mod_plsql` are as follows:

- You can invoke PL/SQL Web applications such as Application Express without installing Oracle HTTP Server, thereby simplifying installation, configuration, and administration of PL/SQL based Web applications.
- You use the same configuration approach that is used to deliver content from Oracle XML DB in response to FTP and HTTP requests.

PL/SQL Web Toolkit

This set of PL/SQL packages is a generic interface that enables you to use stored subprograms invoked by `mod_plsql` at run time.

In response to a browser request, a PL/SQL subprogram updates or retrieves data from Oracle Database according to the user input. It then generates an HTTP response to the browser, typically in the form of a file download or HTML to be displayed. The Web Toolkit API enables stored subprograms to perform actions such as:

- Obtain information about an HTTP request
- Generate HTTP headers such as content-type and mime-type
- Set browser cookies
- Generate HTML pages

Table 9–1 describes commonly used PL/SQL Web Toolkit packages.

Table 9–1 Commonly Used Packages in the PL/SQL Web Toolkit

Package	Description of Contents
`HTF`	Function versions of the subprograms in the `htp` package. The function versions do not directly generate output in a Web page. Instead, they pass their output as return values to the statements that invoke them. Use these functions when you must nest function calls.
`HTP`	Subprograms that generate HTML tags. For example, the procedure `htp.anchor` generates the HTML anchor tag, `<A>`.
`OWA_CACHE`	Subprograms that enable the PL/SQL gateway cache feature to improve performance of your PL/SQL Web application. You can use this package to enable expires-based and validation-based caching with the PL/SQL gateway file system.
`OWA_COOKIE`	Subprograms that send and retrieve HTTP cookies to and from a client Web browser. Cookies are strings a browser uses to maintain state between HTTP calls. State can be maintained throughout a client session or longer if a cookie expiration date is included.
`OWA_CUSTOM`	The authorize function used by cookies.
`OWA_IMAGE`	Subprograms that obtain the coordinates where a user clicked an image. Use this package when you have an image map whose destination links invoke a PL/SQL gateway.

Table 9–1 (Cont.) Commonly Used Packages in the PL/SQL Web Toolkit

Package	Description of Contents
OWA_OPT_LOCK	Subprograms that impose database optimistic locking strategies to prevent lost updates. Lost updates can otherwise occur if a user selects, and then attempts to update, a row whose values were changed in the meantime by another user.
OWA_PATTERN	Subprograms that perform string matching and string manipulation with regular expressions.
OWA_SEC	Subprograms used by the PL/SQL gateway for authenticating requests.
OWA_TEXT	Subprograms used by package OWA_PATTERN for manipulating strings. You can also use them directly.
OWA_UTIL	These types of utility subprograms: ▪ Dynamic SQL utilities to produce pages with dynamically generated SQL code. ▪ HTML utilities to retrieve the values of CGI environment variables and perform URL redirects. ▪ Date utilities for correct date-handling. Date values are simple strings in HTML, but must be properly treated as an Oracle Database data type.
WPG_DOCLOAD	Subprograms that download documents from a document repository that you define using the DAD configuration.

See Also: *Oracle Database PL/SQL Packages and Types Reference* for syntax, descriptions, and examples for the PL/SQL Web Toolkit packages

Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application

As explained in detail in the *Oracle HTTP Server mod_plsql User's Guide*, `mod_plsql` maps Web client requests to PL/SQL stored subprograms over HTTP. See this documentation for instructions.

See Also:

- *Oracle HTTP Server mod_plsql User's Guide* to learn how to configure and use `mod_plsql`
- *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server* for information about the `mod_plsql` module

Using Embedded PL/SQL Gateway

The embedded gateway functions very similar to the `mod_plsql` gateway. Before using the embedded version of the gateway, familiarize yourself with the *Oracle HTTP Server mod_plsql User's Guide*. Much of the information is the same or similar.

Topics:

- How Embedded PL/SQL Gateway Processes Client Requests
- Installing Embedded PL/SQL Gateway
- Configuring Embedded PL/SQL Gateway
- Invoking PL/SQL Stored Subprograms Through Embedded PL/SQL Gateway

- Securing Application Access with Embedded PL/SQL Gateway
- Restrictions in Embedded PL/SQL Gateway
- Using Embedded PL/SQL Gateway: Scenario

How Embedded PL/SQL Gateway Processes Client Requests

Figure 9–2 illustrates the process by which the embedded gateway handles client HTTP requests.

Figure 9–2 Processing Client Requests with Embedded PL/SQL Gateway

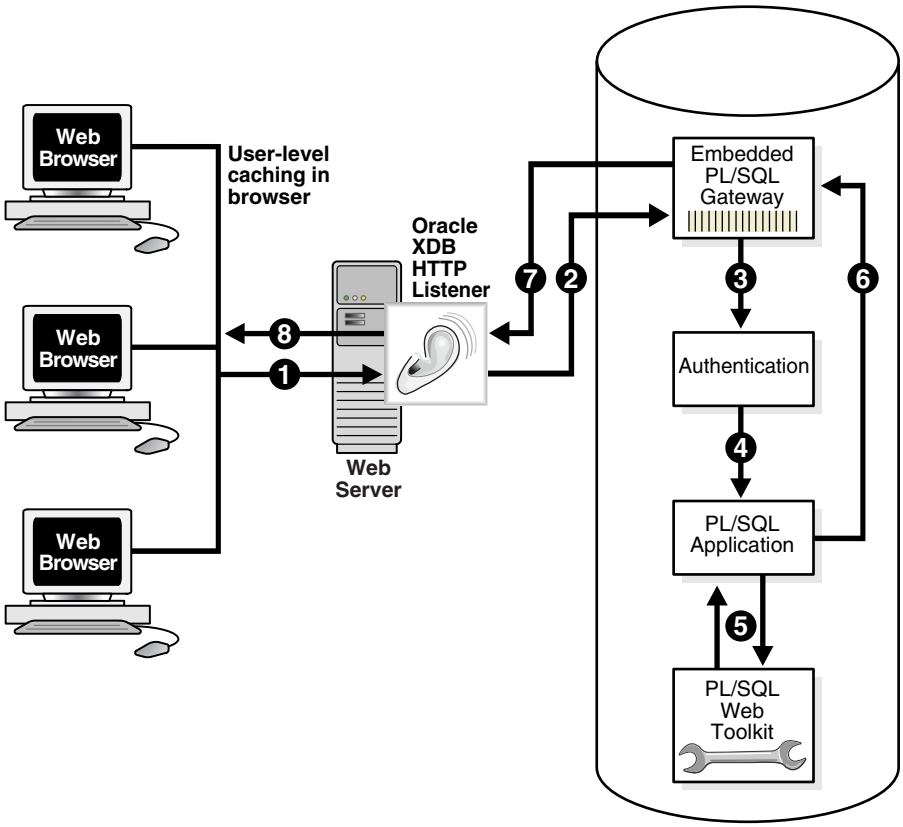

The explanation of the steps in Figure 9–2 is as follows:

1. The Oracle XML DB HTTP Listener receives a request from a client browser to request to invoke a PL/SQL subprogram. The subprogram can either be written directly in PL/SQL or indirectly generated when a PL/SQL Server Page is uploaded to the database and compiled.
2. The XML DB HTTP Listener routes the request to the embedded PL/SQL gateway as specified in its virtual-path mapping configuration.
3. The embedded gateway uses the HTTP request information and the gateway configuration to determine which database account to use for authentication.
4. The embedded gateway prepares the call parameters and invokes the PL/SQL subprogram in the application.
5. The PL/SQL subprogram generates an HTML page out of relational data and the PL/SQL Web Toolkit accessed from the database.

6. The application sends the page to the embedded gateway.
7. The embedded gateway sends the page to the XML DB HTTP Listener.
8. The XML DB HTTP Listener sends the page to the client browser.

Unlike `mod_plsql`, the embedded gateway processes HTTP requests with the Oracle XML DB Listener. This listener is the same server-side process as the Oracle Net Listener and supports Oracle Net Services, HTTP, and FTP.

Configure general HTTP listener settings through the XML DB interface (for instructions, see *Oracle XML DB Developer's Guide*). Configure the HTTP listener either by using Oracle Enterprise Manager or by editing the `xdbconfig.xml` file. Use the `DBMS_EPG` package for all embedded PL/SQL gateway configuration, for example, creating or setting attributes for a DAD.

Installing Embedded PL/SQL Gateway

The embedded gateway requires these components:

- XML DB HTTP Listener
- PL/SQL Web Toolkit

The embedded PL/SQL gateway is installed as part of Oracle XML DB. If you are using a preconfigured database created during an installation or by the Database Configuration Assistant (DBCA), then Oracle XML DB is installed and configured. For information about manually adding Oracle XML DB to an existing database, see *Oracle XML DB Developer's Guide*.

The PL/SQL Web Toolkit is part of the standard installation of the database, so no supplementary installation is necessary.

Configuring Embedded PL/SQL Gateway

You configure `mod_plsql` by editing the Oracle HTTP Server configuration files. Because the embedded gateway is installed as part of the Oracle XML DB HTTP Listener, you manage the embedded gateway as a servlet through the Oracle XML DB servlet management interface.

The configuration interface to the embedded gateway is the PL/SQL package `DBMS_EPG`. This package modifies the underlying `xdbconfig.xml` configuration file that XML DB uses. The default values of the embedded gateway configuration parameters are sufficient for most users.

Topics:

- Configuring Embedded PL/SQL Gateway: Overview
- Configuring User Authentication for Embedded PL/SQL Gateway

Configuring Embedded PL/SQL Gateway: Overview

As in `mod_plsql`, each request for a PL/SQL stored subprogram is associated with a **Database Access Descriptor (DAD)**. A DAD is a set of configuration values used for database access. A DAD specifies information such as:

- The database account to use for authentication
- The subprogram to use for uploading and downloading documents

In the embedded PL/SQL gateway, a DAD is represented as a servlet in the XML DB HTTP Listener configuration. Each DAD attribute maps to an XML element in the

configuration file `xdbconfig.xml`. The value of the DAD attribute corresponds to the element content. For example, the `database-username` DAD attribute corresponds to the `<database-username>` XML element; if the value of the DAD attribute is `HR` it corresponds to `<database-username>HR<database-username>`. DAD attribute names are case-sensitive.

Use the `DBMS_EPG` package to perform these embedded PL/SQL gateway configurations:

1. Create a DAD with the `DBMS_EPG.CREATE_DAD` procedure.
2. Set DAD attributes with the `DBMS_EPG.SET_DAD_ATTRIBUTE` procedure.

   All DAD attributes are optional. If you do not specify an attribute, it has its initial value.

Table 9–2 lists the embedded PL/SQL gateway attributes and the corresponding `mod_plsql` DAD parameters. Enumeration values in the "Legal Values" column are case-sensitive.

Table 9–2 Mapping Between mod_plsql and Embedded PL/SQL Gateway DAD Attributes

mod_plsql DAD Attribute	Embedded PL/SQL Gateway DAD Attribute	Multiple Occurr.	Legal Values
`PlsqlAfterProcedure`	`after-procedure`	No	String
`PlsqlAlwaysDescribeProcedure`	`always-describe-procedure`	No	Enumeration of On, Off
`PlsqlAuthenticationMode`	`authentication-mode`	No	Enumeration of Basic, SingleSignOn, GlobalOwa, CustomOwa, PerPackageOwa
`PlsqlBeforeProcedure`	`before-procedure`	No	String
`PlsqlBindBucketLengths`	`bind-bucket-lengths`	Yes	Unsigned integer
`PlsqlBindBucketWidths`	`bind-bucket-widths`	Yes	Unsigned integer
`PlsqlCGIEnvironmentList`	`cgi-environment-list`	Yes	String
`PlsqlCompatibilityMode`	`compatibility-mode`	No	Unsigned integer
`PlsqlDatabaseUsername`	`database-username`	No	String
`PlsqlDefaultPage`	`default-page`	No	String
`PlsqlDocumentPath`	`document-path`	No	String
`PlsqlDocumentProcedure`	`document-procedure`	No	String
`PlsqlDocumentTablename`	`document-table-name`	No	String
`PlsqlErrorStyle`	`error-style`	No	Enumeration of ApacheStyle, ModplsqlStyle, DebugStyle
`PlsqlExclusionList`	`exclusion-list`	Yes	String
`PlsqlFetchBufferSize`	`fetch-buffer-size`	No	Unsigned integer
`PlsqlInfoLogging`	`info-logging`	No	Enumeration of InfoDebug
`PlsqlInputFilterEnable`	`input-filter-enable`	No	String
`PlsqlMaxRequestsPerSession`	`max-requests-per-session`	No	Unsigned integer
`PlsqlNLSLanguage`	`nls-language`	No	String
`PlsqlOWADebugEnable`	`owa-debug-enable`	No	Enumeration of On, Off
`PlsqlPathAlias`	`path-alias`	No	String
`PlsqlPathAliasProcedure`	`path-alias-procedure`	No	String
`PlsqlRequestValidationFunction`	`request-validation-function`	No	String
`PlsqlSessionCookieName`	`session-cookie-name`	No	String

Table 9–2 (Cont.) Mapping Between mod_plsql and Embedded PL/SQL Gateway DAD Attributes

mod_plsql DAD Attribute	Embedded PL/SQL Gateway DAD Attribute	Multiple Occurr.	Legal Values
`PlsqlSessionStateManagement`	`session-state-management`	No	Enumeration of StatelessWithResetPackageState, StatelessWithFastRestPackageState, StatelessWithPreservePackageState
`PlsqlTransferMode`	`transfer-mode`	No	Enumeration of Char, Raw
`PlsqlUploadAsLongRaw`	`upload-as-long-raw`	No	String

The default values of the DAD attributes are sufficient for most users of the embedded gateway. `mod_plsql` users do not need these attributes:

- `PlsqlDatabasePassword`
- `PlsqlDatabaseConnectString` (because the embedded gateway does not support logon to external databases)

Like the DAD attributes, the global configuration parameters are optional. Table 9–3 describes the `DBMS_EPG` global attributes and the corresponding `mod_plsql` global parameters.

Table 9–3 Mapping Between mod_plsql and Embedded PL/SQL Gateway Global Attributes

mod_plsql DAD Attribute	Embedded PL/SQL Gateway DAD Attribute	Multiple Occurr.	Legal Values
`PlsqlLogLevel`	`log-level`	No	Unsigned integer
`PlsqlMaxParameters`	`max-parameters`	No	Unsigned integer

See Also:

- *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server* for detailed descriptions of the `mod_plsql` DAD attributes. See this documentation for default values and usage notes.
- *Oracle Database PL/SQL Packages and Types Reference* to learn about the `DBMS_EPG` package
- *Oracle XML DB Developer's Guide* for an account of the `xdbconfig.xml` file

Configuring User Authentication for Embedded PL/SQL Gateway

Because it uses the XML DB authentication schemes, the embedded gateway handles database authentication differently from `mod_plsql`. In particular, it does not store database passwords in a DAD.

Note: To serve a PL/SQL Web application on the Internet but maintain the database behind a firewall, do not use the embedded PL/SQL gateway to run the application; use `mod_plsql`.

Use the `DBMS_EPG` package to configure database authentication.

Topics:

- Configuring Static Authentication with DBMS_EPG
- Configuring Dynamic Authentication with DBMS_EPG

- Configuring Anonymous Authentication with DBMS_EPG
- Determining the Authentication Mode of a DAD
- Creating and Configuring DADs: Examples
- Determining the Authentication Mode for a DAD: Example
- Determining the Authentication Mode for All DADs: Example
- Showing DAD Authorizations that Are Not in Effect: Example
- Examining Embedded PL/SQL Gateway Configuration

Configuring Static Authentication with DBMS_EPG Static authentication is for the `mod_plsql` user who stores database user names and passwords in the DAD so that the browser user is not required to enter database authentication information.

To configure static authentication, follow these steps:

1. Log on to the database as an XML DB administrator (that is, a user with the `XDBADMIN` role assigned).
2. Create the DAD. For example, this procedure creates a DAD invoked `HR_DAD` and maps the virtual path to `/hrweb/`:

   ```
   EXEC DBMS_EPG.CREATE_DAD('HR_DAD', '/hrweb/*');
   ```

3. For this step, you need the `ALTER ANY USER` system privilege. Set the DAD attribute `database-username` to the database account whose privileges must be used by the DAD. For example, this procedure specifies that the DAD named `HR_DAD` has the privileges of the `HR` account:

   ```
   EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('HR_DAD', 'database-username', 'HR');
   ```

   The DAD attribute `database-username` is case-sensitive.
4. Assign the DAD the privileges of the database user specified in the previous step. This authorization enables end users to invoke procedures and access document tables through the embedded PL/SQL gateway with the privileges of the authorized account. For example:

   ```
   EXEC DBMS_EPG.AUTHORIZE_DAD('HR_DAD', 'HR');
   ```

   Alternatively, you can log off as the user with `XDBADMIN` privileges, log on as the database user whose privileges must be used by the DAD, and then use this command to assign these privileges to the DAD:

   ```
   EXEC DBMS_EPG.AUTHORIZE_DAD('HR_DAD');
   ```

   > **Note:** Multiple users can authorize the same DAD. The `database-username` attribute setting of the DAD determines which user's privileges to use.

Unlike `mod_plsql`, the embedded gateway connects to the database as the special user `ANONYMOUS`, but accesses database objects with the user privileges assigned to the DAD. The database rejects access if the browser user attempts to connect explicitly with the HTTP `Authorization` header.

> **Note:** The account ANONYMOUS is locked after XML DB installation. To use static authentication with the embedded PL/SQL gateway, first unlock this account.

Configuring Dynamic Authentication with DBMS_EPG Dynamic authentication is for the mod_plsql user who does not store database user names and passwords in the DAD.

In dynamic authentication, a database user does not have to authorize the embedded gateway to use its privileges to access database objects. Instead, browser users must supply the database authentication information through the HTTP Basic Authentication scheme.

The action of the embedded gateway depends on whether the database-username attribute is set for the DAD. If the attribute is not set, then the embedded gateway connects to the database as the user supplied by the browser client. If the attribute is set, then the database restricts access to the user specified in the database-username attribute.

To set up dynamic authentication, follow these steps:

1. Log on to the database as a an XML DB administrator (that is, a user with the XDBADMIN role).
2. Create the DAD. For example, this procedure creates a DAD invoked DYNAMIC_DAD and maps the virtual path to /hrweb/:

   ```
   EXEC DBMS_EPG.CREATE_DAD('DYNAMIC_DAD', '/hrweb/*');
   ```

3. Optionally, set the DAD attribute database-username to the database account whose privileges must be used by the DAD. The browser prompts the user to enter the username and password for this account when accessing the DAD. For example, this procedure specifies that the DAD named DYNAMIC_DAD has the privileges of the HR account:

   ```
   EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('DYNAMIC_DAD', 'database-username', 'HR');
   ```

   The attribute database-username is case-sensitive.

> **WARNING: Passwords sent through the HTTP Basic Authentication scheme are not encrypted. Configure the embedded gateway to use the HTTPS protocol to protect the passwords sent by the browser clients.**

Configuring Anonymous Authentication with DBMS_EPG Anonymous authentication is for the mod_plsql user who creates a special DAD database user for database logon, but stores the application procedures and document tables in a different schema and grants access to the procedures and document tables to PUBLIC.

To set up anonymous authentication, follow these steps:

1. Log on to the database as an XML DB administrator, that is, a user with the XDBADMIN role assigned.
2. Create the DAD. For example, this procedure creates a DAD invoked HR_DAD and maps the virtual path to /hrweb/:

   ```
   EXEC DBMS_EPG.CREATE_DAD('HR_DAD', '/hrweb/*');
   ```

3. Set the DAD attribute `database-username` to `ANONYMOUS`. For example:

```
EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('HR_DAD', 'database-username', 'ANONYMOUS');
```

   Both `database-username` and `ANONYMOUS` are case-sensitive.

   You need not authorize the embedded gateway to use `ANONYMOUS` privileges to access database objects, because `ANONYMOUS` has no system privileges and owns no database objects.

Determining the Authentication Mode of a DAD If you know the name of a DAD, then the authentication mode for this DAD depends on these factors:

- Does the DAD exist?
- Is the `database-username` attribute for the DAD set?
- Is the DAD authorized to use the privilege of the `database-username` user?
- Is the `database-username` attribute the one that the user authorized to use the DAD?

Table 9–4 shows how the answers to the preceding questions determine the authentication mode.

Table 9–4 Authentication Possibilities for a DAD

DAD Exists?	database-username set?	User authorized?	Mode
Yes	Yes	Yes	Static
Yes	Yes	No	Dynamic restricted
Yes	No	Does not matter	Dynamic
Yes	Yes (to `ANONYMOUS`)	Does not matter	Anonymous
No			N/A

For example, assume that you create a DAD named `MY_DAD`. If the `database-username` attribute for `MY_DAD` is set to `HR`, but the `HR` user does not authorize `MY_DAD`, then the authentication mode for `MY_DAD` is dynamic and restricted. A browser user who attempts to run a PL/SQL subprogram through `MY_DAD` is prompted to enter the `HR` database username and password.

The `DBA_EPG_DAD_AUTHORIZATION` view shows which users have authorized use of a DAD. The `DAD_NAME` column displays the name of the DAD; the `USERNAME` column displays the user whose privileges are assigned to the DAD. The DAD authorized might not exist.

> **See Also:** *Oracle Database Reference* for more information about the `DBA_EPG_DAD_AUTHORIZATION` view

Creating and Configuring DADs: Examples Example 9–1 does this:

- Creates a DAD with static authentication for database user `HR` and assigns it the privileges of the `HR` account, which then authorizes it.
- Creates a DAD with dynamic authentication that is not restricted to any user.
- Creates a DAD with dynamic authentication that is restricted to the `HR` account.

Example 9–1 Creating and Configuring DADs

```
-----------------------------------------------------------------------
```

```
--- DAD with static authentication
-----------------------------------------------------------------------

CONNECT SYSTEM AS SYSDBA
PASSWORD: password
EXEC DBMS_EPG.CREATE_DAD('Static_Auth_DAD', '/static/*');
EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('Static_Auth_DAD', 'database-username', 'HR');
GRANT EXECUTE ON DBMS_EPG TO HR;

-- Authorization
CONNECT HR
PASSWORD: password
EXEC DBMS_EPG.AUTHORIZE_DAD('Static_Auth_DAD');

------------------------------------------------------------------------
-- DAD with dynamic authentication
-----------------------------------------------------------------------

CONNECT SYSTEM AS SYSDBA
PASSWORD: password
EXEC DBMS_EPG.CREATE_DAD('Dynamic_Auth_DAD', '/dynamic/*');

-------------------------------------------------------------------------
-- DAD with dynamic authentication restricted
-------------------------------------------------------------------------

EXEC DBMS_EPG.CREATE_DAD('Dynamic_Auth_DAD_Restricted', '/dynamic/*');
EXEC DBMS_EPG.SET_DAD_ATTRIBUTE
  ('Dynamic_Auth_DAD_Restricted', 'database-username', 'HR');
```

The creation and authorization of a DAD are independent; therefore you can:

- Authorize a DAD that does not exist (it can be created later)
- Authorize a DAD for which you are not the user (however, the authorization does not take effect until the DAD `database-user` attribute is changed to your username)

Example 9–2 creates a DAD with static authentication for database user `HR` and assigns it the privileges of the `HR` account. Then:

- Instead of authorizing that DAD, the database user `HR` authorizes a nonexistent DAD.

  Although the user might have done this by mistake, no error occurs, because the nonexistent DAD might be created later.

- The database user `OE` authorizes the DAD (whose `database-user` attribute is set to `HR`.

  No error occurs, but the authorization does not take effect until the DAD `database-user` attribute is changed to `OE`.

Example 9–2 Authorizing DADs to be Created or Changed Later

```
REM Create DAD with static authentication for database user HR

CONNECT SYSTEM AS SYSDBA
PASSWORD: password
EXEC DBMS_EPG.CREATE_DAD('Static_Auth_DAD', '/static/*');
EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('Static_Auth_DAD', 'database-username', 'HR');
GRANT EXECUTE ON DBMS_EPG TO HR;
```

```
REM Database user HR authorizes DAD that does not exist

CONNECT HR
PASSWORD: password
EXEC DBMS_EPG.AUTHORIZE_DAD('Static_Auth_DAD_Typo');

REM Database user OE authorizes DAD with database-username 'HR'

CONNECT OE
PASSWORD: password
EXEC DBMS_EPG.AUTHORIZE_DAD('Static_Auth_DAD');
```

Determining the Authentication Mode for a DAD: Example Example 9–3 creates a PL/SQL procedure, `show_dad_auth_status`, which accepts the name of a DAD and reports its authentication mode. If the specified DAD does not exist, the procedure exits with an error.

Example 9–3 Determining the Authentication Mode for a DAD

```
CREATE OR REPLACE PROCEDURE show_dad_auth_status (p_dadname VARCHAR2) IS
  v_daduser VARCHAR2(32);
  v_cnt     PLS_INTEGER;
BEGIN
  -- Determine DAD user
  v_daduser := DBMS_EPG.GET_DAD_ATTRIBUTE(p_dadname, 'database-username');

  -- Determine whether DAD authorization exists for DAD user
  SELECT COUNT(*)
    INTO v_cnt
      FROM DBA_EPG_DAD_AUTHORIZATION da
        WHERE da.DAD_NAME = p_dadname
          AND da.USERNAME = v_daduser;

  -- If DAD authorization exists for DAD user, authentication mode is static
  IF (v_cnt > 0) THEN
    DBMS_OUTPUT.PUT_LINE (
      '''' || p_dadname ||
      ''' is set up for static authentication for user ''' ||
      v_daduser || '''.');
    RETURN;
  END IF;

  -- If no DAD authorization exists for DAD user, authentication mode is dynamic

  -- Determine whether dynamic authentication is restricted to particular user
  IF (v_daduser IS NOT NULL) THEN
    DBMS_OUTPUT.PUT_LINE (
      '''' || p_dadname ||
      ''' is set up for dynamic authentication for user ''' ||
      v_daduser || ''' only.');
  ELSE
    DBMS_OUTPUT.PUT_LINE (
      '''' || p_dadname ||
      ''' is set up for dynamic authentication for any user.');
  END IF;
END;
/
```

Assume that you have run the script in Example 9–1 to create and configure various DADs. The output is:

```
SET SERVEROUTPUT ON;
BEGIN
  show_dad_auth_status('Static_Auth_DAD');
END;
/
'Static_Auth_DAD' is set up for static authentication for user 'HR'.
```

Determining the Authentication Mode for All DADs: Example The anonymous block in Example 9–4 reports the authentication modes of all registered DADs. It invokes the `show_dad_auth_status` procedure from Example 9–3.

Example 9–4 Showing the Authentication Mode for All DADs

```
DECLARE
  v_dad_names DBMS_EPG.VARCHAR2_TABLE;
BEGIN
  DBMS_OUTPUT.PUT_LINE
    ('---------- Authorization Status for All DADs ----------');
  DBMS_EPG.GET_DAD_LIST(v_dad_names);

  FOR i IN 1..v_dad_names.count LOOP
    show_dad_auth_status(v_dad_names(i));
  END LOOP;
END;
/
```

If you have run the script in Example 9–1 to create and configure various DADs, the output of Example 9–4 is:

```
---------- Authorization Status for All DADs ----------
'Static_Auth_DAD' is set up for static auth for user 'HR'.
'Dynamic_Auth_DAD' is set up for dynamic auth for any user.
'Dynamic_Auth_DAD_Restricted' is set up for dynamic auth for user 'HR' only.
```

Showing DAD Authorizations that Are Not in Effect: Example The anonymous block in Example 9–5 reports DAD authorizations that are *not* in effect. A DAD authorization is not in effect in either of these situations:

- The user who authorizes the DAD is not the user specified by the `database-username` attribute of the DAD
- The user authorizes a DAD that does not exist

Example 9–5 Showing DAD Authorizations that Are Not in Effect

```
DECLARE
  v_dad_names DBMS_EPG.VARCHAR2_TABLE;
  v_dad_user  VARCHAR2(32);
  v_dad_found BOOLEAN;
BEGIN
  DBMS_OUTPUT.PUT_LINE
    ('---------- DAD Authorizations Not in Effect ----------');
  DBMS_EPG.GET_DAD_LIST(v_dad_names);

  FOR r IN (SELECT * FROM DBA_EPG_DAD_AUTHORIZATION) LOOP  -- Outer loop
    v_dad_found := FALSE;
    FOR i IN 1..v_dad_names.count LOOP  -- Inner loop
      IF (r.DAD_NAME = v_dad_names(i)) THEN
```

```
        v_dad_user :=
          DBMS_EPG.GET_DAD_ATTRIBUTE(r.DAD_NAME, 'database-username');

        -- Is database-username the user for whom DAD is authorized?
        IF (r.USERNAME <> v_dad_user) THEN
          DBMS_OUTPUT.PUT_LINE (
            'DAD authorization of ''' || r.dad_name ||
            ''' by user ''' || r.username || '''' ||
            ' is not in effect because DAD user is ' ||
            '''' || v_dad_user || '''.');
        END IF;
        v_dad_found := TRUE;
        EXIT;  -- Inner loop
      END IF;
    END LOOP;  -- Inner loop

    -- Does DAD exist?
    IF (NOT v_dad_found) THEN
      DBMS_OUTPUT.PUT_LINE (
       'DAD authorization of ''' || r.dad_name ||
       ''' by user ''' || r.username ||
       ''' is not in effect because the DAD does not exist.');
    END IF;
  END LOOP;  -- Outer loop
END;
/
```

If you have run the script in Example 9–2 to create and configure various DADs, the output of Example 9–5 (reformatted to fit on the page) is:

```
---------- DAD Authorizations Not in Effect ----------
DAD authorization of 'Static_Auth_DAD' by user 'OE' is not in effect
  because DAD user is 'HR'.
DAD authorization of 'Static_Auth_DAD_Typo' by user 'HR' is not in effect
  because DAD does not exist.
```

Examining Embedded PL/SQL Gateway Configuration When you are connected to the database as a user with system privileges, this script helps you examine the configuration of the embedded PL/SQL gateway:

```
$ORACLE_HOME/rdbms/admin/epgstat.sql
```

Example 9–6 shows the output of the `epgstat.sql` script for Example 9–1 when the `ANONYMOUS` account is locked.

Example 9–6 epgstat.sql Script Output for Example 9–1

Command to run script:

```
@$ORACLE_HOME/rdbms/admin/epgstat.sql
```

Result:

```
+-------------------------------------+
| XDB protocol ports:                 |
|  XDB is listening for the protocol  |
|  when the protocol port is nonzero. |
+-------------------------------------+

HTTP Port FTP Port
--------- --------
```

```
        0        0

1 row selected.

+---------------------------+
| DAD virtual-path mappings |
+---------------------------+

Virtual Path                     DAD Name
-------------------------------- --------------------------------
/dynamic/*                       Dynamic_Auth_DAD_Restricted
/static/*                        Static_Auth_DAD

2 rows selected.

+----------------+
| DAD attributes |
+----------------+

DAD Name        DAD Param             DAD Value
------------    --------------------- ----------------------------------------
Dynamic_Auth    database-username     HR
_DAD_Restric
ted

Static_Auth_    database-username     HR
DAD


2 rows selected.

+----------------------------------------------------+
| DAD authorization:                                 |
|  To use static authentication of a user in a DAD, |
|  the DAD must be authorized for the user.          |
+----------------------------------------------------+

DAD Name                         User Name
-------------------------------- --------------------------------
Static_Auth_DAD                  HR
                                 OE
Static_Auth_DAD_Typo             HR

3 rows selected.

+----------------------------+
| DAD authentication schemes |
+----------------------------+

DAD Name             User Name                        Auth Scheme
-------------------- -------------------------------- ------------------
Dynamic_Auth_DAD                                      Dynamic
Dynamic_Auth_DAD_Res HR                               Dynamic Restricted
tricted

Static_Auth_DAD      HR                               Static

3 rows selected.

+-------------------------------------------------------+
```

```
| ANONYMOUS user status:                                  |
|  To use static or anonymous authentication in any DAD, |
|  the ANONYMOUS account must be unlocked.                |
+-------------------------------------------------------+

Database User   Status
--------------- --------------------
ANONYMOUS       EXPIRED & LOCKED

1 row selected.

+-----------------------------------------------------------------+
| ANONYMOUS access to XDB repository:                             |
|  To allow public access to XDB repository without authentication, |
|  ANONYMOUS access to the repository must be allowed.            |
+-----------------------------------------------------------------+

Allow repository anonymous access?
----------------------------------
false

1 row selected.
```

Invoking PL/SQL Stored Subprograms Through Embedded PL/SQL Gateway

The basic steps for invoking PL/SQL subprograms through the embedded PL/SQL gateway are the same as for the `mod_plsql` gateway. See *Oracle HTTP Server mod_plsql User's Guide* for instructions. You must adapt the `mod_plsql` instructions slightly for use with the embedded gateway. For example, invoke the embedded gateway in a browser by entering the URL in this format:

```
protocol://hostname[:port]/virt-path/[[!][schema.][package.]proc_name[?query_str]]
```

The placeholder `virt-path` stands for the virtual path that you configured in `DBMS_EPG.CREATE_DAD`. The `mod_plsql` documentation uses `DAD_location` instead of `virt-path`.

These topics documented in *Oracle HTTP Server mod_plsql User's Guide* apply equally to the embedded gateway:

- Transaction mode
- Supported data types
- Parameter-passing scheme
- File upload and download support
- Path-aliasing
- Common Gateway Interface (CGI) environment variables

Securing Application Access with Embedded PL/SQL Gateway

The embedded gateway shares the same protection mechanism with `mod_plsql`. See *Oracle HTTP Server mod_plsql User's Guide* for instructions.

Restrictions in Embedded PL/SQL Gateway

The `mod_plsql` restrictions documented in the first chapter of *Oracle HTTP Server mod_plsql User's Guide* apply equally to the embedded gateway. In addition, the embedded version of the gateway does not support these features:

- Dynamic HTML caching
- System monitoring
- Authentication modes other than Basic

  For information about authentication modes, see *Oracle HTTP Server mod_plsql User's Guide*.

Using Embedded PL/SQL Gateway: Scenario

This section illustrates how to write a simple application that queries the `hr.employees` table and delivers HTML output to a Web browser through the PL/SQL gateway. It assumes that you have both XML DB and the sample schemas installed.

To write and run the program follow these steps:

1. Log on to the database as a user with `ALTER USER` privileges and ensure that the database account `ANONYMOUS` is unlocked. The `ANONYMOUS` account, which is locked by default, is required for static authentication. If the account is locked, then use this SQL statement to unlock it:

   ```
   ALTER USER anonymous ACCOUNT UNLOCK;
   ```

2. Log on to the database as an XML DB administrator, that is, a user with the `XDBADMIN` role.

   To determine which users and roles were granted the `XDADMIN` role, query the data dictionary:

   ```
   SELECT *
   FROM DBA_ROLE_PRIVS
   WHERE GRANTED_ROLE = 'XDBADMIN';
   ```

3. Create the DAD. For example, this procedure creates a DAD invoked `HR_DAD` and maps the virtual path to `/plsql/`:

   ```
   EXEC DBMS_EPG.CREATE_DAD('HR_DAD', '/plsql/*');
   ```

4. Set the DAD attribute `database-username` to the database user whose privileges must be used by the DAD. For example, this procedure specifies that the DAD `HR_DAD` accesses database objects with the privileges of user `HR`:

   ```
   EXEC DBMS_EPG.SET_DAD_ATTRIBUTE('HR_DAD', 'database-username', 'HR');
   ```

   The attribute `database-username` is case-sensitive.

5. Grant `EXECUTE` privilege to the database user whose privileges must be used by the DAD (so that he or she can authorize the DAD). For example:

   ```
   GRANT EXECUTE ON DBMS_EPG TO HR;
   ```

6. Log off as the XML DB administrator and log on to the database as the database user whose privileges must be used by the DAD (for example, `HR`).

7. Authorize the embedded PL/SQL gateway to invoke procedures and access document tables through the DAD. For example:

```
EXEC DBMS_EPG.AUTHORIZE_DAD('HR_DAD');
```

8. Create a sample PL/SQL stored procedure invoked `print_employees`. This program creates an HTML page that includes the result set of a query of `hr.employees`:

```
CREATE OR REPLACE PROCEDURE print_employees IS
  CURSOR emp_cursor IS
    SELECT last_name, first_name
      FROM hr.employees
        ORDER BY last_name;
BEGIN
  HTP.PRINT('<html>');
  HTP.PRINT('<head>');
  HTP.PRINT('<meta http-equiv="Content-Type" content="text/html">');
  HTP.PRINT('<title>List of Employees</title>');
  HTP.PRINT('</head>');
  HTP.PRINT('<body TEXT="#000000" BGCOLOR="#FFFFFF">');
  HTP.PRINT('<h1>List of Employees</h1>');
  HTP.PRINT('<table width="40%" border="1">');
  HTP.PRINT('<tr>');
  HTP.PRINT('<th align="left">Last Name</th>');
  HTP.PRINT('<th align="left">First Name</th>');
  HTP.PRINT('</tr>');
  FOR emp_record IN emp_cursor LOOP
    HTP.PRINT('<tr>');
    HTP.PRINT('<td>' || emp_record.last_name  || '</td>');
    HTP.PRINT('<td>' || emp_record.first_name || '</td>');
  END LOOP;
  HTP.PRINT('</table>');
  HTP.PRINT('</body>');
  HTP.PRINT('</html>');
END;
/
```

9. Ensure that the Oracle Net listener can accept HTTP requests. You can determine the status of the listener on Linux and UNIX by running this command at the system prompt:

```
lsnrctl status | grep HTTP
```

Output (reformatted from a single line to multiple lines from page size constraints):

```
(DESCRIPTION=
  (ADDRESS=(PROTOCOL=tcp)(HOST=example.com)(PORT=8080))
  (Presentation=HTTP)
  (Session=RAW)
)
```

If you do not see the HTTP service started, then you can add these lines to your initialization parameter file (replacing *`listener_name`* with the name of your Oracle Net local listener), then restart the database and the listener:

```
dispatchers="(PROTOCOL=TCP)"
local_listener=listener_name
```

10. Run the `print_employees` program from your Web browser. For example, you can use this URL, replacing *`host`* with the name of your host computer and *`port`* with the value of the `PORT` parameter in the previous step:

```
http://host:port/plsql/print_employees
```

For example, if your host is `test.com` and your HTTP port is `8080`, then enter:

```
http://example.com:8080/plsql/print_employees
```

The Web browser returns an HTML page with a table that includes the first and last name of every employee in the `hr.employees` table.

Generating HTML Output with PL/SQL

Traditionally, PL/SQL Web applications use function calls to generate each HTML tag for output. These functions are part of the PL/SQL Web Toolkit packages that come with Oracle Database. Example 9–7 shows how to generate a simple HTML page by calling the `HTP` functions that correspond to each HTML tag.

Example 9–7 Using HTP Functions to Generate HTML Tags

```
CREATE OR REPLACE PROCEDURE html_page IS
BEGIN
  HTP.HTMLOPEN;                            -- generates <HTML>
  HTP.HEADOPEN;                            -- generates <HEAD>
  HTP.TITLE('Title');                      -- generates <TITLE>Hello</TITLE>
  HTP.HEADCLOSE;                           -- generates </HTML>

  -- generates <BODY TEXT="#000000" BGCOLOR="#FFFFFF">
  HTP.BODYOPEN( cattributes => 'TEXT="#000000" BGCOLOR="#FFFFFF"');

  -- generates <H1>Heading in the HTML File</H1>
  HTP.HEADER(1, 'Heading in the HTML File');

  HTP.PARA;                                 -- generates <P>
  HTP.PRINT('Some text in the HTML file.');
  HTP.BODYCLOSE;                            -- generates </BODY>
  HTP.HTMLCLOSE;                            -- generates </HTML>
END;
/
```

An alternative to making function calls that correspond to each tag is to use the `HTP.PRINT` function to print both text and tags. Example 9–8 illustrates this technique.

Example 9–8 Using HTP.PRINT to Generate HTML Tags

```
CREATE OR REPLACE PROCEDURE html_page2 IS
BEGIN
  HTP.PRINT('<html>');
  HTP.PRINT('<head>');
  HTP.PRINT('<meta http-equiv="Content-Type" content="text/html">');
  HTP.PRINT('<title>Title of the HTML File</title>');
  HTP.PRINT('</head>');
  HTP.PRINT('<body TEXT="#000000" BGCOLOR="#FFFFFF">');
  HTP.PRINT('<h1>Heading in the HTML File</h1>');
  HTP.PRINT('<p>Some text in the HTML file.');
  HTP.PRINT('</body>');
  HTP.PRINT('</html>');
END;
/
```

Chapter 10, "Developing PL/SQL Server Pages (PSP)," describes an additional method for delivering using PL/SQL to generate HTML content. PL/SQL server pages enables you to build on your knowledge of HTML tags and avoid learning a new set of function calls. In an application written as a set of PL/SQL server pages, you can still use functions from the PL/SQL Web toolkit to:

- Simplify the processing involved in displaying tables
- Store persistent data (cookies)
- Work with CGI protocol internals

Passing Parameters to PL/SQL Web Applications

To be useful in a wide variety of situations, a Web application must be interactive enough to allow user choices. To keep the attention of impatient Web surfers, streamline the interaction so that users can specify these choices very simply, without excessive decision-making or data entry.

The main methods of passing parameters to PL/SQL Web applications are:

- Using HTML form tags. The user fills in a form on one Web page, and all the data and choices are transmitted to a stored subprogram when the user clicks the `Submit` button on the page.
- Hard-coded in the URL. The user clicks on a link, and a set of predefined parameters are transmitted to a stored subprogram. Typically, you include separate links on your Web page for all the choices that the user might want.

Topics:

- Passing List and Dropdown-List Parameters from an HTML Form
- Passing Option and Check Box Parameters from an HTML Form
- Passing Entry-Field Parameters from an HTML Form
- Passing Hidden Parameters from an HTML Form
- Uploading a File from an HTML Form
- Submitting a Completed HTML Form
- Handling Missing Input from an HTML Form
- Maintaining State Information Between Web Pages

Passing List and Dropdown-List Parameters from an HTML Form

List boxes and drop-down lists are implemented with the HTML tag `<SELECT>`.

Use a list box for a large number of choices or to allow multiple selections. List boxes are good for showing items in alphabetical order so that users can find an item quickly without reading all the choices.

Use a drop-down list in these situations:

- There are a small number of choices
- Screen space is limited.
- Choices are in an unusual order.

The drop-down captures the attention of first-time users and makes them read the items. If you keep the choices and order consistent, then users can memorize the

motion of selecting an item from the drop-down list, allowing them to make selections quickly as they gain experience. Example 9–9 shows a simple drop-down list.

Example 9–9 HTML Drop-Down List

```
<form>
<select name="seasons">
<option value="winter">Winter
<option value="spring">Spring
<option value="summer">Summer
<option value="fall">Fall
</select>
```

Passing Option and Check Box Parameters from an HTML Form

Options pass either a null value (if none of the options in a group is checked), or the value specified on the option that is checked.

To specify a default value for a set of options, you can include the `CHECKED` attribute in an`INPUT` tag, or include a `DEFAULT` clause on the parameter within the stored subprogram. When setting up a group of options, be sure to include a choice that indicates "no preference", because after selecting a option, the user can select a different one, but cannot clear the selection completely. For example, include a "Don't Care" or "Don't Know" selection along with "Yes" and "No" choices, in case someone makes a selection and then realizes it was wrong.

Check boxes need special handling, because your stored subprogram might receive a null value, a single value, or multiple values:

All the check boxes with the same `NAME` attribute comprise a check box group. If none of the check boxes in a group is checked, the stored subprogram receives a null value for the corresponding parameter.

If one check box in a group is checked, the stored subprogram receives a single `VARCHAR2` parameter.

If multiple check boxes in a group are checked, the stored subprogram receives a parameter with the PL/SQL type `TABLE OF VARCHAR2`. You must declare a type like `TABLE OF VARCHAR2`, or use a predefined one like `OWA_UTIL.IDENT_ARR`. To retrieve the values, use a loop:

```
CREATE OR REPLACE PROCEDURE handle_checkboxes (
  checkboxes owa_util.ident_arr
) AS
BEGIN
  FOR i IN 1..checkboxes.count
  LOOP
    htp.print('<p>Check Box value: ' || checkboxes(i));
  END LOOP;
END;
/
```

Passing Entry-Field Parameters from an HTML Form

Entry fields require the most validation, because a user might enter data in the wrong format, out of range, and so on. If possible, validate the data on the client side using a client-side Javascript function, and format it correctly for the user or prompt them to enter it again.

For example:

- You might prevent the user from entering alphabetic characters in a numeric entry field, or from entering characters after reaching a length limit.
- You might silently remove spaces and dashes from a credit card number if the stored subprogram expects the value in that format.
- You might inform the user immediately when they type a number that is too large, so that they can retype it.

Because you cannot always rely on such validation to succeed, code the stored subprograms to deal with these cases anyway. Rather than forcing the user to use the `Back` button when they enter wrong data, display a single page with an error message and the original form with all the other values filled in.

For sensitive information such as passwords, a special form of the entry field, `<INPUT TYPE=PASSWORD>`, hides the text as it is typed in.

The procedure in Example 9–10 accepts two strings as input. The first time the procedure is invoked, the user sees a simple form prompting for the input values. When the user submits the information, the same procedure is invoked again to check if the input is correct. If the input is OK, the procedure processes it. If not, the procedure prompts for input, filling in the original values for the user.

Example 9–10 Passing Entry-Field Parameters from an HTML Form

```
DROP TABLE name_zip_table;
CREATE TABLE name_zip_table (
  name     VARCHAR2(100),
  zipcode  NUMBER
);

-- Store a name and associated zip code in the database.
CREATE OR REPLACE PROCEDURE associate_name_with_zipcode
  (name VARCHAR2 := NULL,
   zip  VARCHAR2 := NULL)
AS
BEGIN
  -- Each entry field must contain a value. Zip code must be 6 characters.
  -- (In a real program you perform more extensive checking.)

  IF name IS NOT NULL AND zip IS NOT NULL AND length(zip) = 6 THEN
    INSERT INTO name_zip_table (name, zipcode) VALUES (name, zip);
    HTP.PRINT('<p>The person ' || name || ' has the zip code ' || zip || '.');

    -- If input was OK, stop here. User does not see form again.
    RETURN;
  END IF;

  -- If user entered incomplete or incorrect data, show error message.

  IF (name IS NULL AND zip IS NOT NULL)
    OR (name IS NOT NULL AND zip IS NULL)
      OR (zip IS NOT NULL AND length(zip) != 6)
  THEN
    HTP.PRINT('<p><b>Please reenter data. Fill all fields,
              and use 6-digit zip code.</b>');
  END IF;

  -- If user entered no data or incorrect data, show error message
  -- & make form invoke same procedure to check input values.
```

```
  HTP.FORMOPEN('HR.associate_name_with_zipcode', 'GET');
  HTP.PRINT('<p>Enter your name:</td>');
  HTP.PRINT
   ('<td valign=center><input type=text name=name value="' || name || '">');
  HTP.PRINT('<p>Enter your zip code:</td>');
  HTP.PRINT
    ('<td valign=center><input type=text name=zip value="' || zip || '">');
  HTP.FORMSUBMIT(NULL, 'Submit');
  HTP.FORMCLOSE;
END;
/
```

Passing Hidden Parameters from an HTML Form

One technique for passing information through a sequence of stored subprograms, without requiring the user to specify the same choices each time, is to include hidden parameters in the form that invokes a stored subprogram. The first stored subprogram places information, such as a user name, into the HTML form that it generates. The value of the hidden parameter is passed to the next stored subprogram, as if the user had entered it through a option or entry field.

Other techniques for passing information from one stored subprogram to another include:

- Sending a "cookie" containing the persistent information to the browser. The browser then sends this same information back to the server when accessing other Web pages from the same site. Cookies are set and retrieved through the HTTP headers that are transferred between the browser and the Web server before the HTML text of each Web page.
- Storing the information in the database itself, where later stored subprograms can retrieve it. This technique involves some extra overhead on the database server, and you must still find a way to keep track of each user as multiple users access the server at the same time.

Uploading a File from an HTML Form

You can use an HTML form to choose a file on a client system, and transfer it to the server. A stored subprogram can insert the file into the database as a `CLOB`, `BLOB`, or other type that can hold large amounts of data.

The PL/SQL Web toolkit and the PL/SQL gateway have the notion of a "document table" that holds uploaded files.

> **See Also:** *mod_plsql User's Guide*

Submitting a Completed HTML Form

By default, an HTML form must have a `Submit` button, which transmits the data from the form to a stored subprogram or CGI program. You can label this button with text of your choice, such as "Search", "Register", and so on.

You can have multiple forms on the same page, each with its own form elements and `Submit` button. You can even have forms consisting entirely of hidden parameters, where the user makes no choice other than clicking the button.

Using JavaScript or other scripting languages, you can eliminate the Submit button and have the form submitted in response to some other action, such as selecting from a drop-down list. This technique is best when the user only makes a single selection, and the confirmation step of the `Submit` button is not essential.

Handling Missing Input from an HTML Form

When an HTML form is submitted, your stored subprogram receives null parameters for any form elements that are not filled in. For example, null parameters can result from an empty entry field, a set of check boxes, options, or list items with none checked, or a `VALUE` parameter of "" (empty quotation marks).

Regardless of any validation you do on the client side, always code stored subprograms to handle the possibility that some parameters are null:

- Specify an initial value in all parameter declarations, to prevent an exception when the stored subprogram is invoked with a missing form parameter. You can set the initial value to zero for numeric values (when that makes sense), and to `NULL` when you want to check whether the user actually specifies a value.
- Before using an input parameter value that has the initial value `NULL`, check if it is null.
- Make the subprogram generate sensible results even when not all input parameters are specified. You might leave some sections out of a report, or display a text string or image in a report to indicate where parameters were not specified.
- Provide a way to fill in the missing values and run the stored subprogram again, directly from the results page. For example, include a link that invokes the same stored subprogram with an additional parameter, or display the original form with its values filled in as part of the output.

Maintaining State Information Between Web Pages

Web applications are particularly concerned with the idea of **state**, the set of data that is current at a particular moment in time. It is easy to lose state information when switching from one Web page to another, which might result in asking the user to make the same choices over and over.

You can pass state information between dynamic Web pages using HTML forms. The information is passed as a set of name-value pairs, which are turned into stored subprogram parameters for you.

If the user has to make multiple selections, or one selection from many choices, or it is important to avoid an accidental selection, use an HTML form. After the user makes and reviews all the choices, they confirm the choices with the `Submit` button. Subsequent pages can use forms with hidden parameters (`<INPUT TYPE=HIDDEN>` tags) to pass these choices from one page to the next.

If the user is only considering one or two choices, or the decision points are scattered throughout the Web page, you can save the user from hunting around for the `Submit` button by representing actions as hyperlinks and including any necessary name-value pairs in the query string (the part after the `?` within a URL).

An alternative way to main state information is to use Oracle Application Server and its `mod_ose` module. This approach lets you store state information in package variables that remain available as a user moves around a Web site.

> **See Also:** The Oracle Application Server documentation set at http://www.oracle.com/technology/documentation

Performing Network Operations in PL/SQL Subprograms

Oracle provides packages that allow PL/SQL subprograms to perform these network operations:

- Sending E-Mail from PL/SQL
- Getting a Host Name or Address from PL/SQL
- Using TCP/IP Connections from PL/SQL
- Retrieving HTTP URL Contents from PL/SQL
- Using Tables, Image Maps, Cookies, and CGI Variables from PL/SQL

Internet Protocol version 6 (IPv6) Support

As of Oracle Database 11*g* Release 2 (11.2.0.1), PL/SQL network utility packages support IPv6 addresses. The package interfaces have not changed: Any interface parameter that expects a network host accepts an IPv6 address in string form, and any interface that returns an IP address can return an IPv6 address.

However, applications that use network addresses might need small changes, and recompilation, to accommodate IPv6 addresses. An IPv6 address has 128 bits, while an IPv4 address has only 32 bits. In a URL, an IPv6 address must be enclosed in brackets. For example:

```
http://[2001:0db8:85a3:08d3:1319:8a2e:0370:7344]/
```

See Also:

- *Oracle Database Net Services Administrator's Guide* for detailed information about IPv6 support in Oracle Database
- *Oracle Database PL/SQL Packages and Types Reference* for information about IPv6 support in specific PL/SQL network utility packages

Sending E-Mail from PL/SQL

Using the `UTL_SMTP` package, a PL/SQL subprogram can send e-mail, as in Example 9–11.

Example 9–11 Sending E-Mail from PL/SQL

```
CREATE OR REPLACE PROCEDURE send_test_message
IS
  mailhost   VARCHAR2(64) := 'mailhost.example.com';
  sender     VARCHAR2(64) := 'me@example.com';
  recipient  VARCHAR2(64) := 'you@example.com';
  mail_conn  UTL_SMTP.CONNECTION;
BEGIN
  mail_conn := UTL_SMTP.OPEN_CONNECTION(mailhost, 25); -- 25 is the port
  UTL_SMTP.HELO(mail_conn, mailhost);
  UTL_SMTP.MAIL(mail_conn, sender);
  UTL_SMTP.RCPT(mail_conn, recipient);

  UTL_SMTP.OPEN_DATA(mail_conn);
  UTL_SMTP.WRITE_DATA(mail_conn, 'This is a test message.' || chr(13));
  UTL_SMTP.WRITE_DATA(mail_conn, 'This is line 2.' || chr(13));
  UTL_SMTP.CLOSE_DATA(mail_conn);

  /* If message were in single string, open_data(), write_data(),
     and close_data() could be in a single call to data(). */

  UTL_SMTP.QUIT(mail_conn);
EXCEPTION
  WHEN OTHERS THEN
```

```
  -- Insert error-handling code here
  NULL;
END;
/
```

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for detailed information about the UTL_SMTP package

Getting a Host Name or Address from PL/SQL

Using the UTL_INADDR package, a PL/SQL subprogram can determine the host name of the local system or the IP address of a given host name.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for detailed information about the UTL_INADDR package

Using TCP/IP Connections from PL/SQL

Using the UTL_TCP package, a PL/SQL subprogram can open TCP/IP connections to systems on the network, and read or write to the corresponding sockets.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for detailed information about the UTL_TCP package

Retrieving HTTP URL Contents from PL/SQL

Using the UTL_HTTP package, a PL/SQL subprogram can:

- Retrieve the contents of an HTTP URL

  The contents are usually in the form of HTML-tagged text, but might be any kind of file that can be downloaded from a Web server (for example, plain text or a JPEG image).
- Control HTTP session details (such as headers, cookies, redirects, proxy servers, IDs and passwords for protected sites, and CGI parameters)
- Speed up multiple accesses to the same Web site, using HTTP 1.1 persistent connections

A PL/SQL subprogram can construct and interpret URLs for use with the UTL_HTTP package by using the functions UTL_URL.ESCAPE and UTL_URL.UNESCAPE.

The PL/SQL procedure in Example 9–12 uses the UTL_HTTP package to retrieve the contents of an HTTP URL.

Example 9–12 Retrieving HTTP URL Contents from PL/SQL

```
CREATE OR REPLACE PROCEDURE show_url
  (url      IN VARCHAR2,
  username IN VARCHAR2 := NULL,
  password IN VARCHAR2 := NULL)
AS
  req        UTL_HTTP.REQ;
  resp       UTL_HTTP.RESP;
  name_      VARCHAR2(256);
  value_     VARCHAR2(1024);
  data_      VARCHAR2(255);
  my_scheme  VARCHAR2(256);
  my_realm   VARCHAR2(256);
  my_proxy   BOOLEAN;
```

```
BEGIN
  -- When going through a firewall, pass requests through this host.
  -- Specify sites inside the firewall that do not need the proxy host.

  UTL_HTTP.SET_PROXY('proxy.example.com', 'corp.example.com');

  -- Ask UTL_HTTP not to raise an exception for 4xx and 5xx status codes,
  -- rather than just returning the text of the error page.

  UTL_HTTP.SET_RESPONSE_ERROR_CHECK(FALSE);

  -- Begin retrieving this Web page.
 req := UTL_HTTP.BEGIN_REQUEST(url);

  -- Identify yourself.
  --  Some sites serve special pages for particular browsers.
  UTL_HTTP.SET_HEADER(req, 'User-Agent', 'Mozilla/4.0');

  -- Specify user ID and password for pages that require them.
  IF (username IS NOT NULL) THEN
    UTL_HTTP.SET_AUTHENTICATION(req, username, password);
  END IF;

  -- Start receiving the HTML text.
  resp := UTL_HTTP.GET_RESPONSE(req);

  -- Show status codes and reason phrase of response.
  DBMS_OUTPUT.PUT_LINE('HTTP response status code: ' || resp.status_code);
  DBMS_OUTPUT.PUT_LINE
    ('HTTP response reason phrase: ' || resp.reason_phrase);

  -- Look for client-side error and report it.
  IF (resp.status_code >= 400) AND (resp.status_code <= 499) THEN

  -- Detect whether page is password protected
  -- and you didn't supply the right authorization.

  IF (resp.status_code = UTL_HTTP.HTTP_UNAUTHORIZED) THEN
  UTL_HTTP.GET_AUTHENTICATION(resp, my_scheme, my_realm, my_proxy);
  IF (my_proxy) THEN
    DBMS_OUTPUT.PUT_LINE('Web proxy server is protected.');
    DBMS_OUTPUT.PUT('Please supply the required ' || my_scheme ||
      ' authentication username for realm ' || my_realm ||
      ' for the proxy server.');
  ELSE
    DBMS_OUTPUT.PUT_LINE('Web page ' || url || ' is protected.');
    DBMS_OUTPUT.PUT('Please supplied the required ' || my_scheme ||
      ' authentication username for realm ' || my_realm ||
      ' for the Web page.');
  END IF;
  ELSE
    DBMS_OUTPUT.PUT_LINE('Check the URL.');
  END IF;

  UTL_HTTP.END_RESPONSE(resp);
    RETURN;

  -- Look for server-side error and report it.
  ELSIF (resp.status_code >= 500) AND (resp.status_code <= 599) THEN
    DBMS_OUTPUT.PUT_LINE('Check if the Web site is up.');
```

```
    UTL_HTTP.END_RESPONSE(resp);
    RETURN;
  END IF;

  -- HTTP header lines contain information about cookies, character sets,
  -- and other data that client and server can use to customize each
  -- session.

  FOR i IN 1..UTL_HTTP.GET_HEADER_COUNT(resp) LOOP
    UTL_HTTP.GET_HEADER(resp, i, name_, value_);
    DBMS_OUTPUT.PUT_LINE(name_ || ': ' || value_);
  END LOOP;

  -- Read lines until none are left and an exception is raised.
  LOOP
    UTL_HTTP.READ_LINE(resp, value_);
    DBMS_OUTPUT.PUT_LINE(value_);
  END LOOP;
EXCEPTION
  WHEN UTL_HTTP.END_OF_BODY THEN
  UTL_HTTP.END_RESPONSE(resp);
END;
/
```

This block shows examples of calls to the procedure in Example 9–12, but the URLs are for nonexistent pages. Substitute URLs from your own Web server.

```
BEGIN
  show_url('http://www.oracle.com/no-such-page.html');
  show_url('http://www.oracle.com/protected-page.html');
  show_url
    ('http://www.oracle.com/protected-page.html','username','password');
END;
/
```

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for detailed information about the `UTL_HTTP` package
- *Oracle Database PL/SQL Packages and Types Reference* for detailed information about `UTL_URL.ESCAPE` and `UTL_URL.UNESCAPE`

Using Tables, Image Maps, Cookies, and CGI Variables from PL/SQL

Using packages supplied by Oracle, and the `mod_plsql` plug-in of Oracle HTTP Server (OHS), a PL/SQL subprogram can format the results of a query in an HTML table, produce an image map, set and get HTTP cookies, check the values of CGI variables, and perform other typical Web operations.

Documentation for these packages is not part of the database documentation library. The location of the documentation depends on your application server. To get started with these packages, look at their subprogram names and parameters using the SQL*Plus `DESCRIBE` statement:

```
DESCRIBE HTP;
DESCRIBE HTF;
DESCRIBE OWA_UTIL;
```

10

Developing PL/SQL Server Pages (PSP)

This chapter explains how to develop PL/SQL Server Pages (PSP), which let you include dynamic content in web pages.

Topics:

- What Are PL/SQL Server Pages and Why Use Them?
- Prerequisites for Developing and Deploying PL/SQL Server Pages
- PL/SQL Server Pages and the HTP Package
- PL/SQL Server Pages and Other Scripting Solutions
- Developing PL/SQL Server Pages
- Loading PL/SQL Server Pages into the Database
- Querying PL/SQL Server Page Source Code
- Running PL/SQL Server Pages Through URLs
- Examples of PL/SQL Server Pages
- Debugging PL/SQL Server Pages
- Putting PL/SQL Server Pages into Production

What Are PL/SQL Server Pages and Why Use Them?

PL/SQL Server Pages (PSP) are server-side scripts that include dynamic content, including the results of SQL queries, inside web pages. You can author the web pages in an HTML authoring tool and insert blocks of PL/SQL code.

Example 10–1 shows a simple PL/SQL server page called `simple.psp`.

Example 10–1 simple.psp

```
<%@ page language="PL/SQL" %>
<%@ page contentType="text/html" %>
<%@ plsql procedure="show_employees" %>
<%-- This example displays the last name and first name of every
     employee in the hr.employees table. --%>
<%!
  CURSOR emp_cursor IS
    SELECT last_name, first_name
      FROM hr.employees
        ORDER BY last_name;
%>
<html>
```

```
<head>
<meta http-equiv="Content-Type" content="text/html">
<title>List of Employees</title>
</head>
<body TEXT="#000000" BGCOLOR="#FFFFFF">
<h1>List of Employees</h1>
<table width="40%" border="1">
<tr>
<th align="left">Last Name</th>
<th align="left">First Name</th>
</tr>
<%  FOR emp_record IN emp_cursor LOOP %>
  <tr>
  <td> <%= emp_record.last_name %> </td>
  <td> <%= emp_record.first_name %> </td>
  </tr>
<%  END LOOP; %>
</table>
</body>
</html>
```

You can compile and load a PL/SQL server page into the database with the `loadpsp` command-line utility. This command loads `simple.psp` into the `hr` schema, replacing the `show_employees` procedure if it exists:

```
loadpsp -replace simple.psp
Enter Password: password
```

Browser users can run the `show_employees` procedure through a URL. An HTML page that displays the last and first names of employees in the `hr.employees` table is returned to the browser through the PL/SQL gateway.

Deploying content through PL/SQL Server Pages has these advantages:

- For developers familiar with PL/SQL, the server pages are the easiest way to create professional web pages that include database-generated content. You can develop web pages as you usually do and then embed PL/SQL code in the HTML.
- PL/SQL Server Pages can be more convenient than using the `HTP` and `HTF` packages to write out HTML content line by line.
- Because processing is performed on the database server, the client browser receives a plain HTML page with no special script tags. You can support all browsers and browser levels equally.
- Network traffic is efficient because use of PL/SQL Server Pages minimizes the number of database round-trips.
- You can write content quickly and follow a rapid, iterative development process. You maintain central control of the software, with only a web browser required on the client system.

Prerequisites for Developing and Deploying PL/SQL Server Pages

To develop and deploy PL/SQL server pages, you must meet these prerequisites:

- To write a PL/SQL server page you need access to a text editor or HTML authoring tool for writing the script. No other development tool is required.
- To load a PL/SQL server page you need:
  - An account on the database in which to load the server pages.

  - Execution rights to the `loadpsp` command-line utility, which is located in `$ORACLE_HOME/bin`.

- To deploy the server pages you must use mod_plsql. As explained in "Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application" on page 9-4, the gateway uses the PL/SQL Web Toolkit.

  See Also:

  - "Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application" on page 9-4

PL/SQL Server Pages and the HTP Package

You can enable browser users to run PL/SQL units through HTTP in these ways:

- By writing an HTML page with embedded PL/SQL code and compiling it as a PL/SQL server page. You might invoke subprograms from the PL/SQL Web Toolkit, but not to generate every line of HTML output.
- By writing a complete stored subprogram that produces HTML by invoking the `HTP` and `OWA_*` packages in the PL/SQL Web Toolkit. For information about this technique, see "Generating HTML Output with PL/SQL" on page 9-20.

Thus, you must choose which technique to use when writing your web application. The key factors in choosing between these techniques are:

- What source are you using as a starting point?
  - If you have a large body of HTML, and want to include dynamic content or make it the front end of a database application, then use PL/SQL Server Pages.
  - If you have a large body of PL/SQL code that produces formatted output, then you might find it more convenient to produce HTML tags by changing your print statements to invoke the `HTP` package of the PL/SQL Web Toolkit.
- What is the fastest and most convenient authoring environment for your group?
  - If most work is done using HTML authoring tools, then use PL/SQL Server Pages.
  - If you use authoring tools that produce PL/SQL code, then it might be less convenient to use PL/SQL Server Pages.

PL/SQL Server Pages and Other Scripting Solutions

Scripting solutions can be client-side or server-side. JavaScript is a very popular client-side scripting languages. PL/SQL Server Pages fully support JavaScript. Because any kind of tags can be passed unchanged to the browser through a PL/SQL server page, you can include JavaScript or other client-side script code in a PL/SQL server page.

Java Server Pages (JSP) and Active Server Pages (ASP) are two of the most popular server-side scripting solutions. Compared to PL/SQL Server Pages:

- Java server pages are loosely analogous to PL/SQL Server Pages pages; Java servlets are analogous to PL/SQL packages. PL/SQL Server Pages use the same script tag syntax as JSP to make it easy to switch back and forth.
- PL/SQL Server Pages use syntax that is similar to ASP, although not identical. Typically, you must translate from VBScript or JScript to PL/SQL. The best

candidates for migration are pages that use the Active Data Object (ADO) interface to perform database operations.

Note: You cannot mix PL/SQL server pages with other server-side script features, such as server-side includes. Often, you can get the same results by using the corresponding PL/SQL Server Pages features.

Developing PL/SQL Server Pages

To develop a PL/SQL server page, you can start with an existing web page or with an existing stored subprogram. Either way, with a few additions and changes you can create dynamic web pages that perform database operations and display the results.

The file for a PL/SQL server page must have the extension `.psp`. It can contain whatever content you choose, with text and tags interspersed with PL/SQL Server Pages directives, declarations, and scriptlets. A server page can take these forms:

- In the simplest case, it is an HTML file. Compiling it as a PL/SQL server page produces a stored subprogram that outputs the same HTML file.
- In the most complex case, it is a PL/SQL subprogram that generates all the content of the web page, including the tags for title, body, and headings.
- In the typical case, it is a mixture of HTML (providing the static parts of the page) and PL/SQL (providing the dynamic content).

The order and placement of the PL/SQL Server Pages directives and declarations is usually not significant. It becomes significant only when another file is included. For ease of maintenance, Oracle recommends that you put the directives and declarations near the beginning of the file.

Table 10–1 lists the PL/SQL Server Pages elements and directs you to the section that explains how to use them. The section "Using Quotation Marks and Escaping Strings in a PSP Script" on page 10-12 describes how to use quotation marks in strings that are used in various PL/SQL Server Pages elements.

Table 10–1 PSP Elements

PSP Element	Name	Specifies . . .	Section
`<%@ page ... %>`	Page Directive	Characteristics of the PL/SQL server page.	"Specifying Basic Server Page Characteristics" on page 10-5
`<%@ parameter ... %>`	Parameter Directive	The name, and optionally the type and default, for each parameter expected by the PSP stored procedure.	"Accepting User Input" on page 10-8
`<%@ plsql ... %>`	Procedure Directive	The name of the stored procedure produced by the PSP file.	"Naming the PL/SQL Stored Procedure" on page 10-9
`<%@ include ... %>`	Include Directive	The name of a file to be included at a specific point in the PSP file.	"Including the Contents of Other Files" on page 10-9
`<%! ... %>`	Declaration Block	The declaration for a set of PL/SQL variables that are visible throughout the page, not just within the next `BEGIN/END` block.	"Declaring Global Variables in a PSP Script" on page 10-10

Table 10–1 (Cont.) PSP Elements

PSP Element	Name	Specifies . . .	Section
`<% ... %>`	Code Block	A set of PL/SQL statements to be executed when the procedure is run.	"Specifying Executable Statements in a PSP Script" on page 10-10
`<%= ... %>`	Expression Block	A single PL/SQL expression, such as a string, arithmetic expression, function call, or combination of these.	"Substituting Expression Values in a PSP Script" on page 10-11
`<%-- ... --%>`	Comment	A comment in a PSP script.	"Including Comments in a PSP Script" on page 10-12

Note: If you are familiar with dynamic HTML, you can go directly to "Examples of PL/SQL Server Pages" on page 10-16.

Topics:

- Specifying Basic Server Page Characteristics
- Accepting User Input
- Naming the PL/SQL Stored Procedure
- Including the Contents of Other Files
- Declaring Global Variables in a PSP Script
- Specifying Executable Statements in a PSP Script
- Substituting Expression Values in a PSP Script
- Using Quotation Marks and Escaping Strings in a PSP Script
- Including Comments in a PSP Script

Specifying Basic Server Page Characteristics

Use the `<%@ page ... %>` directive to specify characteristics of the PL/SQL server page such as:

- What scripting language it uses.
- What type of information (MIME type) it produces.
- What code to run to handle all uncaught exceptions. This might be an HTML file with a friendly message, renamed to a `.psp` file. You must specify this same file name in the `loadpsp` command that compiles the main PSP file. You must specify the same name in both the `errorPage` directive and in the `loadpsp` command, including any relative path name such as `../include/`.

This code shows the syntax of the `page` directive (the attribute names `contentType` and `errorPage` are case-sensitive):

```
<%@ page
language='PL/SQL'
contentType='content_type_string'
charset='encoding'
errorPage='file.psp'
%>
```

Topics:

- Specifying the Scripting Language

- Returning Data to the Client Browser
- Handling Script Errors

Specifying the Scripting Language

To identify a file as a PL/SQL server page, include this directive somewhere in the file:

```
<%@ page language="PL/SQL" %>
```

This directive is for compatibility with other scripting environments. Example 10–1 shows an example of a simple PL/SQL server page that includes the language directive.

Returning Data to the Client Browser

Options:

- Returning HTML
- Returning XML, Text, and Other Document Types
- Returning Pages Containing Different Character Sets

Returning HTML The PL/SQL parts of a PL/SQL server page are enclosed within special delimiters. All other content is passed exactly as it is—including any white space—to the browser. To display text or HTML tags, write it as you would write a typical web page. You need not invoke any output functions. As illustration, the server page in Example 10–1 returns the HTML page shown in Example 10–2, except that it includes the table rows for the queried employees.

Example 10–2 Sample Returned HTML Page

```
<html>
<head>
<meta http-equiv="Content-Type" content="text/html">
<title>List of Employees</title>
</head>
<body TEXT="#000000" BGCOLOR="#FFFFFF">
<h1>List of Employees</h1>
<table width="40%" border="1">
<tr>
<th align="left">Last Name</th>
<th align="left">First Name</th>
</tr>

  <!-- result set of query of hr.employees inserted here -->

</table>
</body>
</html>
```

Sometimes you might want to display one line of output or another, or change the value of an attribute, based on a condition. You can include control structures and variable substitution inside the PSP delimiters, as shown in this code fragment from Example 10–1:

```
<%  FOR emp_record IN emp_cursor LOOP %>
  <tr>
  <td> <%= emp_record.last_name %> </td>
  <td> <%= emp_record.first_name %> </td>
```

```
  </tr>
<%  END LOOP; %>
```

Returning XML, Text, and Other Document Types By default, the PL/SQL gateway transmits files as HTML documents so that the browser interprets the HTML tags. If you want the browser to interpret the document as XML, plain text (with no formatting), or some other document type, then include this directive:

```
<%@ page contentType="MIMEtype" %>
```

The attribute name `contentType` is case-sensitive. Insert `text/html`, `text/xml`, `text/plain`, `image/jpeg`, or some other MIME type that the browser or other client program recognizes. Users might have to configure their browsers to recognize some MIME types. An example of a directive for an Excel spreadsheet is:

```
<%@ page contentType="application/vnd.ms-excel" %>
```

Typically, a PL/SQL server page is intended to be displayed in a web browser. It can also be retrieved and interpreted by a program that can make HTTP requests, such as a a Java or PERL client.

Returning Pages Containing Different Character Sets By default, the PL/SQL gateway transmits files with the character set defined by the PL/SQL gateway. To convert the data to a different character set for browser display, include this directive:

```
<%@ page charset="encoding" %>
```

Specify `Shift_JIS`, `Big5`, `UTF-8`, or another encoding that the client program recognizes.

You must also configure the character set setting in the database accessor descriptor (DAD) of the PL/SQL gateway. Users might have to select the same encoding in their browsers to see the data displayed properly. For example, a database in Japan might have a database character set that uses the `EUC` encoding, but the web browsers are configured to display `Shift_JIS` encoding.

Handling Script Errors

When writing PL/SQL server pages, you can get these types of errors:

- HTML syntax errors

  The browser handles these errors. The `loadpsp` utility does not check for them.

- PL/SQL syntax errors

  The `loadpsp` utility stops and displays the line number, column number, and a brief message. You must fix the error before continuing.

  Any previous version of the stored subprogram can be erased when you attempt to replace it with a script that contains a syntax error. You might want to use one database for prototyping and debugging, and then load the final stored subprogram into a different database for production. You can switch databases using a command-line flag without changing any source code.

- Run-time errors

  To handle database errors that occur when the script runs, you can include PL/SQL exception-handling code within a PSP file and have any unhandled exceptions start a special PL/SQL server page. Use the `errorPage` attribute (the name is case-sensitive) of the `<%@ page ... %>` directive to specify the page name.

The page for unhandled exceptions is a PL/SQL server page with extension .psp. The error subprogram does not receive any parameters, so to determine the cause of the error, it can invoke the SQLCODE and SQLERRM functions. You can also display a standard HTML page without any scripting when an error occurs, but you must still give it the extension .psp and load it into the database as a stored subprogram.

This line specifies errors.psp as the page to run when errors are encountered:

```
<%@ page language="PL/SQL" contentType="text/html" errorPage="errors.psp" %>
```

Accepting User Input

To set up parameter passing for a PL/SQL server page, include a directive with this syntax:

```
<%@ plsql parameter="parameter_name" [type="PL/SQL_type"] [default="value"] %>
```

The default *PL/SQL_type* is VARCHAR2. This directive specifies that the parameter p_employee_id is of the type NUMBER:

```
<%@ plsql parameter="p_employee_id" type="NUMBER" %>
```

Specifying a default value for a parameter makes the parameter optional. The default value is substituted directly into a PL/SQL statement, so any strings must be enclosed in single quotation marks, and you can use special values such as NULL. This directive specifies that the parameter p_last_name has the default value NULL:

```
<%@ plsql parameter="p_last_name" default="NULL" %>
```

User input comes encoded in the URL that retrieves the HTML page. You can generate the URL by hard-coding it in an HTML link, or by invoking your page as the action of an HTML form. Your page receives the input as parameters to a PL/SQL stored subprogram.

Example 10–3 is like Example 10–1, except that it uses a parameter, p_employee_id. If the PL/SQL gateway is configured so that you can run procedures by invoking http://www.host.com/pls/*proc_name*, where *proc_name* is the name of a procedure, then you can pass 200 for parameter p_employee_id as follows:

```
http://www.example.com/pls/show_employees?p_employee_id=200
```

Example 10–3 simplewithuserinput.psp

```
<%@ page language="PL/SQL" %>
<%@ page contentType="text/html" %>
<%@ plsql parameter="p_employee_id" default="null" type="NUMBER" %>
<%@ plsql procedure="show_employees" %>
<%-- This example displays the last name and first name of every
     employee in the hr.employees table. --%>
<%!
  CURSOR emp_cursor IS
  SELECT last_name, first_name
  FROM hr.employees
  WHERE employee_id = p_employee_id
  ORDER BY last_name;
%>
<html>
<head>
<meta http-equiv="Content-Type" content="text/html">
<title>List of Employees</title>
```

```
</head>
<body TEXT="#000000" BGCOLOR="#FFFFFF">
<h1>List of Employees</h1>
<table width="40%" border="1">
<tr>
<th align="left">Last Name</th>
<th align="left">First Name</th>
</tr>
<%  FOR emp_record IN emp_cursor LOOP %>
  <tr>
  <td> <%= emp_record.last_name %> </td>
  <td> <%= emp_record.first_name %> </td>
  </tr>
<%  END LOOP; %>
</table>
</body>
</html>
```

Naming the PL/SQL Stored Procedure

Each top-level PL/SQL server page corresponds to a stored procedure within the server. When you load the page with `loadpsp`, the utility creates a PL/SQL stored procedure. If the server page is *`name`*`.psp`, the default procedure name is *`name`*. For example, if the server page is `hello_world.psp`, then the default procedure name is `hello_world`.

To specify a procedure name, use this directive, where *`procname`* is the name for the procedure:

```
<%@ plsql procedure="procname" %>
```

In Example 10–1, this directive gives the stored procedure the name `show_employees`:

```
<%@ plsql procedure="show_employees" %>
```

It is the name of the procedure, not the name of the PSP script, that you include in the URL.

Including the Contents of Other Files

You can set up an include mechanism to pull in the contents of other files, typically containing either static HTML content or more PL/SQL scripting code. Insert this directive at the point where the content of the other file is to appear, replacing *`filename`* with the name of the file to be included:

```
<%@ include file="filename" %>
```

The included file must have an extension other than `.psp`. You must specify the same name in both the `include` directive and in the `loadpsp` command, including any relative path name such as `../include/`.

Because the files are processed when you load the stored procedure into the database, the substitution is performed only once, not whenever the page is served. Therefore, changes to the included files that occur after the page is loaded into the database are not displayed when the procedure is executed.

You can use the include feature to pull in libraries of code, such as a navigation banners, footers, tables of contents, and so forth into multiple files. Alternatively, you

can use this feature as a macro capability to include the same section of script code in multiple places in a page. This example includes an HTML footer:

```
<%@ include file="footer.htm" %>
```

When you use included files:

- You can use any names and extensions for the included files. For example, you can include a file called `products.txt`.
- If the included files contain PL/SQL scripting code, then they do not need their own set of directives to identify the procedure name, character set, and so on.
- When specifying the names of files to the `loadpsp` utility, you must include the names of all included files also. Specify the names of included files before the names of any `.psp` files.

Declaring Global Variables in a PSP Script

You can use the `<%! ... %>` directive to define a set of PL/SQL variables that are visible throughout the page, not just within the next `BEGIN/END` block. This element typically spans multiple lines, with individual PL/SQL variable declarations ended by semicolons. The syntax for this directive is as follows:

```
<%! PL/SQL declaration;
    [ PL/SQL declaration; ] ... %>
```

The usual PL/SQL syntax is allowed within the block. The delimiters server as shorthand, enabling you to omit the `DECLARE` keyword. All declarations are available to the code later in the file. Example 10–1 includes this cursor declaration:

```
<%!
  CURSOR emp_cursor IS
  SELECT last_name, first_name
  FROM hr.employees
  ORDER BY last_name;
%>
```

You can specify multiple declaration blocks; internally, they are all merged into a single block when the PSP file is created as a stored procedure.

You can also use explicit `DECLARE` blocks within the `<% ... %>` delimiters that are explained in "Specifying Executable Statements in a PSP Script" on page 10-10. These declarations are only visible to the `BEGIN/END` block that follows them.

Note: To make things easier to maintain, keep all your directives and declarations near the beginning of a PL/SQL server page.

Specifying Executable Statements in a PSP Script

You can use the `<% ... %>` code block directive to run a set of PL/SQL statements when the stored procedure is run. This code shows the syntax for executable statements:

```
<% PL/SQL statement;
   [ PL/SQL statement; ] ... %>
```

This element typically spans multiple lines, with individual PL/SQL statements ended by semicolons. The statements can include complete blocks, as in this example, which invokes the `OWA_UTIL.TABLEPRINT` procedure:

```
<% OWA_UTIL.TABLEPRINT(CTABLE => 'hr.employees', CATTRIBUTES => 'border=2',
   CCOLUMNS => 'last_name,first_name', CCLAUSES => 'WHERE employee_id > 100'); %>
```

The statements can also be the bracketing parts of `IF/THEN/ELSE` or `BEGIN/END` blocks. When a code block is split into multiple directives, you can put HTML or other directives in the middle, and the middle pieces are conditionally executed when the stored procedure is run. This code from Example 10–11 provides an illustration of this technique:

```
<% FOR ITEM IN (SELECT product_name, list_price, catalog_url
                FROM product_information
                WHERE list_price IS NOT NULL
                ORDER BY list_price DESC) LOOP
   IF item.list_price > p_minprice THEN
      v_color := '#CCCCFF';
   ELSE
      v_color := '#CCCCCC';
   END IF;
%>
<TR BGCOLOR="<%= v_color %>">
  <TD><A HREF="<%= item.catalog_url %>"><%= item.product_name %></A></TD>
  <TD><BIG><%= item.list_price %></BIG></TD>
</TR>
<% END LOOP; %>
```

All the usual PL/SQL syntax is allowed within the block. The delimiters server as shorthand, letting you omit the `DECLARE` keyword. All the declarations are available to the code later on in the file.

> **Note:** To share procedures, constants, and types across different PL/SQL server pages, compile them into a package in the database by using a plain PL/SQL source file. Although you can reference package procedures, constants, and types from PSP scripts, the PSP scripts can only produce standalone procedures, not packages.

Substituting Expression Values in a PSP Script

An expression directive outputs a single PL/SQL expression, such as a string, arithmetic expression, function call, or combination of these things. The result is substituted as a string at that spot in the HTML page that is produced by the stored procedure. The expression result must be a string value or be able to be cast to a string. For any types that cannot be implicitly cast, such as `DATE`, pass the value to the PL/SQL `TO_CHAR` function.

The syntax of an expression directive is as follows, where the *`expression`* placeholder is replaced by the desired expression:

```
<%= expression %>
```

You need not end the PL/SQL expression with a semicolon.

Example 10–1 includes a directive to print the value of a variable in a row of a cursor:

```
<%= emp_record.last_name %>
```

Compare the preceding example to the equivalent `htp.print` call in this example (note especially the semicolon that ends the statement):

```
<% HTP.PRN (emp_record.last_name); %>
```

The content within the <%= ... %> delimiters is processed by the HTP.PRN function, which trims leading or trailing white space and requires that you enclose literal strings in single quotation marks.

You can use concatenation by using the twin pipe symbol (||) as in PL/SQL. This directive shows an example of concatenation:

```
<%= 'The employee last name is ' || emp_record.last_name %>
```

Using Quotation Marks and Escaping Strings in a PSP Script

PSP attributes use double quotation marks to delimit data. When values specified in PSP attributes are used for PL/SQL operations, they are passed exactly as you specify them in the PSP file. Thus, if PL/SQL requires a string enclosed in single quotation marks, then you must specify the string enclosed in single quotation marks, and enclose the whole thing in double quotation marks.

For example, your PL/SQL procedure might use the string Babe Ruth as the default value for a variable. For the string to be used in PL/SQL, you must enclose it in single quotation marks as 'Babe Ruth'. If you specify this string in the default attribute of a PSP directive, you must enclose it in double quotation marks, like this:

```
<%@ plsql parameter="in_players" default="'Babe Ruth'" %>
```

You can also enclose strings that are enclosed in single quotation marks in another set of single quotation marks. In this case, you must escape the inner single quotation marks by specifying the sequence \'. For example:

```
<%@ plsql parameter="in_players" default="'Walter \'Big Train\' Johnson'" %>
```

You can include most characters and character sequences in a PSP file without having them changed by the PSP loader. To include the sequence %>, specify the escape sequence %\>. To include the sequence <%, specify the escape sequence <\%. For example:

```
<%= 'The %\> sequence is used in scripting language: ' || lang_name %>
<%= 'The <\% sequence is used in scripting language: ' || lang_name %>
```

Including Comments in a PSP Script

To put a comment in the HTML portion of a PL/SQL server page for the benefit of those reading the PSP source code, use this syntax:

```
<%-- PSP comment text --%>
```

Comments in the preceding form do not appear in the HTML output from the PSP and also do not appear when you query the PL/SQL source code in USER_OBJECTS.

To create a comment that is visible in the HTML output and in the USER_OBJECTS source, place the comment in the HTML and use the normal HTML comment syntax:

```
<!-- HTML comment text -->
```

To include a comment inside a PL/SQL block within a PSP, and to make the comment invisible in the HTML output but visible in USER_OBJECTS, use the normal PL/SQL comment syntax, as in this example:

```
-- Comment in PL/SQL code
```

Example 10–4 shows a fragment of a PSP file with the three types of comments.

Example 10–4 Sample Comments in a PSP File

```
<p>Today we introduce our new model XP-10.
<%--
  This is the project with code name "Secret Project".
  Users viewing the HTML page do not see this PSP script comment.
  The comment is not visible in the USER_OBJECTS source code.
--%>
<!--
  Some pictures of the XP-10.
  Users viewing the HTML page source see this comment.
  The comment is also visible in the USER_OBJECTS source code.
-->
<%
FOR image_file IN (SELECT pathname, width, height, description
                   FROM image_library WHERE model_num = 'XP-10')
-- Comments interspersed with PL/SQL statements.
-- Users viewing the HTML page source do not see these PL/SQL comments.
-- These comments are visible in the USER_OBJECTS source code.
LOOP
%>
<img src="<%= image_file.pathname %>" width=<% image_file.width %>
height=<% image_file.height %> alt="<% image_file.description %>">
<br>
<% END LOOP; %>
```

Loading PL/SQL Server Pages into the Database

Use the `loadpsp` utility, which is located in `$ORACLE_HOME/bin`, to load one or more PSP files into the database as stored procedures. Each `.psp` file corresponds to one stored procedure. The pages are compiled and loaded in one step, to speed up the development cycle. The syntax of the `loadpsp` utility is:

```
loadpsp [-replace] [include_file_name...] [error_file_name] psp_file_name...
Enter Password: password
```

When you load a PSP file, the loader performs these actions:

1. Logs on to the database with the specified user name, password, and net service name
2. Creates the stored procedures in the user schema

`-replace` creates procedures with `CREATE OR REPLACE` syntax.

`include_file_name` is the name of a file that is specified in the PSP `include` directive.

`error_file_name` is the name of the file that is specified in the `errorPage` attribute of the PSP `page` directive.

`psp_file_name` is the name of a file that is specified in a PSP `page` directive.

The filenames on the `loadpsp` command line must exactly match the names specified in the PSP `include` and `page` directives, including any relative path name such as `../include/`.

Example 10–5 shows a sample PSP load command.

Example 10–5 Loading PL/SQL Server Pages

```
loadpsp -replace -user joe/abc123@/db3 banner.inc error.psp display_order.psp
```

In Example 10–5:

- The stored procedure is created in the database db3. The database is accessed as user joe with password abc123, both to create the stored procedure and when the stored procedure is executed.
- banner.inc is a file containing boilerplate text and script code that is included by the .psp file. The inclusion occurs when the PSP is loaded into the database, not when the stored procedure is executed.
- error.psp is a file containing code, text, or both that is processed when an unhandled exception occurs, to present a friendly page rather than an internal error message.
- display_order.psp contains the main code and text for the web page. By default, the corresponding stored procedure is named display_order.

Querying PL/SQL Server Page Source Code

The code that loadpsp generates is different from the code in the source file. It has calls to the HTP package, which generates the HTML tags for the web page.

After loading a PSP file, you can see the generated source code by querying the static data dictionary views *_SOURCE. For example, suppose that you load the script in Example 10–1 with this command:

```
loadpsp -replace -user hr simple.psp
Enter Password: password
```

If you log on to the database as user hr, you can view the source code of the PSP as shown in Example 10–6.

Example 10–6 Querying PL/SQL Server Page Source Code

Query:

```
SELECT TEXT
FROM USER_SOURCE
WHERE NAME = 'SHOW_EMPLOYEES'
ORDER BY LINE;
```

Result:

```
PROCEDURE show_employees  AS

  CURSOR emp_cursor IS
  SELECT last_name, first_name
  FROM hr.employees
  ORDER BY last_name;

 BEGIN NULL;
owa_util.mime_header('text/html'); htp.prn('
');
htp.prn('
');
htp.prn('
');
htp.prn('
');
htp.prn('
');
htp.prn('
```

```
<html>
<head>
<meta http-equiv="Content-Type" content="text/html">
<title>List of Employees</title>
</head>
<body TEXT="#000000" BGCOLOR="#FFFFFF">
<h1>List of Employees</h1>
<table width="40%" border="1">
<tr>
<th align="left">Last Name</th>
<th align="left">First Name</th>
</tr>
');
  FOR emp_record IN emp_cursor LOOP
htp.prn('
  <tr>
  <td> ');
htp.prn( emp_record.last_name );
htp.prn(' </td>
  <td> ');
htp.prn( emp_record.first_name );
htp.prn(' </td>
  </tr>
');
  END LOOP;
htp.prn('
</table>
</body>
</html>
');
 END;
```

Running PL/SQL Server Pages Through URLs

After the PL/SQL server page is turned into a stored procedure, you can run the procedure by retrieving an HTTP URL through a web browser or other Internet-aware client program. The virtual path in the URL depends on the way the PL/SQL gateway is configured.

The parameters to the stored procedure are passed through either the `POST` method or the `GET` method of the HTTP protocol. With the `POST` method, the parameters are passed directly from an HTML form and are not visible in the URL. With the `GET` method, the parameters are passed as name-value pairs in the query string of the URL, separated by `&` characters, with most nonalphanumeric characters in encoded format (such as `%20` for a space). You can use the `GET` method to invoke a PSP page from an HTML form, or you can use a hard-coded HTML link to invoke the stored procedure with a given set of parameters.

Using `METHOD=GET`, the syntax of the URL looks something like this:

```
http://sitename/schemaname/procname?parmname1=value1&parmname2=value2
```

For example, this URL includes a `p_lname` and `p_fname` parameter:

```
http://www.example.com/pls/show_employees?p_lname=Ashdown&p_fname=Lance
```

Using `METHOD=POST`, the syntax of the URL does not show the parameters:

```
http://sitename/schemaname/procname
```

For example, this URL specifies a procedure name but does not pass parameters:

```
http://www.example.com/pls/show_employees
```

The METHOD=GET format is more convenient for debugging and allows visitors to pass the same parameters when they return to the page through a bookmark.

The METHOD=POST format allows a larger volume of parameter data, and is suitable for passing sensitive information that must not be displayed in the URL. (URLs linger on in the browser's history list and in the HTTP headers that are passed to the next-visited page.) It is not practical to bookmark pages that are invoked this way.

Examples of PL/SQL Server Pages

This section shows how you might start with a very simple PL/SQL server page, and produce progressively more complicated versions as you gain more confidence.

As you go through each step, you can follow the instructions in "Loading PL/SQL Server Pages into the Database" on page 10-13 and "Running PL/SQL Server Pages Through URLs" on page 10-15 to test the examples.

Topics:

- Setup for PL/SQL Server Pages Examples
- Printing the Sample Table with a Loop
- Allowing a User Selection
- Using an HTML Form to Invoke a PL/SQL Server Page
- Including JavaScript in a PSP File

Setup for PL/SQL Server Pages Examples

These examples use the PRODUCT_INFORMATION table in the OE schema, which is described as follows:

SQL*Plus command:

```
DESCRIBE PRODUCT_INFORMATION;
```

Result:

```
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 PRODUCT_ID                                NOT NULL NUMBER(6)
 PRODUCT_NAME                                       VARCHAR2(50)
 PRODUCT_DESCRIPTION                                VARCHAR2(2000)
 CATEGORY_ID                                        NUMBER(2)
 WEIGHT_CLASS                                       NUMBER(1)
 WARRANTY_PERIOD                                    INTERVAL YEAR(2) TO MONTH
 SUPPLIER_ID                                        NUMBER(6)
 PRODUCT_STATUS                                     VARCHAR2(20)
 LIST_PRICE                                         NUMBER(8,2)
 MIN_PRICE                                          NUMBER(8,2)
 CATALOG_URL                                        VARCHAR2(50)
```

The examples assume:

- You have set up mod_plsql as described in "Using mod_plsql Gateway to Map Client Requests to a PL/SQL Web Application" on page 9-4.

- You have created a DAD for static authentication of the OE user.
- You can access PL/SQL stored procedures created in the OE schema through this URL, where proc_name is the name of a stored procedure:http://www.example.com/pls/proc_name

For debugging purposes, you can display the complete contents of a SQL table with a call to OWA_UTIL.TABLEPRINT, as in Example 10–7. Later examples show other techniques that give more control over the presentation.

Example 10–7 show_prod_simple.psp

```
<%@ plsql procedure="show_prod_simple" %>
<HTML>
<HEAD><TITLE>Show Contents of product_information (Complete Dump)</TITLE></HEAD>
<BODY>
<%
DECLARE
  dummy BOOLEAN;
BEGIN
  dummy := OWA_UTIL.TABLEPRINT('oe.product_information','border');
END;
%>
</BODY>
</HTML>
```

Load the PSP in Example 10–7 at the command line as follows:

```
loadpsp -replace -user oe/password show_prod_simple.psp
Enter Password: password
```

Access the PSP through this URL:

```
http://www.example.com/pls/show_prod_simple
```

Printing the Sample Table with a Loop

Example 10–7 loops through the items in the product_information table and adjusts the SELECT statement to retrieve only a subset of the rows or columns. This example uses a very simple presentation, a set of list items, to avoid any problems from mismatched or unclosed table tags.

Example 10–8 show_catalog_raw.psp

```
<%@ plsql procedure="show_prod_raw" %>
<HTML>
<HEAD><TITLE>Show Products (Raw Form)</TITLE></HEAD>
<BODY>
<UL>
<% FOR item IN (SELECT product_name, list_price, catalog_url
                FROM product_information
                WHERE list_price IS NOT NULL
                ORDER BY list_price DESC) LOOP %>
<LI>
Item = <%= item.product_name %><BR>
Price = <%= item.list_price %><BR>
URL = <%= item.catalog_url %><BR>
<% END LOOP; %>
</UL>
</BODY>
</HTML>
```

Example 10–9 shows a more sophisticated variation of Example 10–8 in which formatting is added to the HTML to improve the presentation.

Example 10–9 show_catalog_pretty.psp

```
<%@ plsql procedure="show_prod_pretty" %>
<HTML>
<HEAD><TITLE>Show Products (Better Form)</TITLE></HEAD>
<BODY>
<UL>
<% FOR item IN (SELECT product_name, list_price, catalog_url
                FROM product_information
                WHERE list_price IS NOT NULL
                ORDER BY list_price DESC) LOOP %>
<LI>
Item = <A HREF=<%= item.catalog_url %>><%= item.product_name %></A><BR>
Price = <BIG><%= item.list_price %></BIG><BR>
<% END LOOP; %>
</UL>
</BODY>
</HTML>
```

Allowing a User Selection

In Example 10–7, Example 10–8, and Example 10–9, the HTML page remains the same unless the `PRODUCT_INFORMATION` table is updated. Example 10–10:

- Makes the HTML page accept a minimum price, and presents only the items that are more expensive. (Your customers' buying criteria might vary.)
- Sets the default minimum price to 100 units of the appropriate currency.

Example 10–10 show_product_partial.psp

```
<%@ plsql procedure="show_product_partial" %>
<%@ plsql parameter="p_minprice" default="100" %>
<HTML>
<HEAD><TITLE>Show Items Greater Than Specified Price</TITLE></HEAD>
<BODY>
<P>This report shows the items whose price is greater than <%= p_minprice %>.
<UL>
<% FOR ITEM IN (SELECT product_name, list_price, catalog_url
                FROM product_information
                WHERE list_price > p_minprice
                ORDER BY list_price DESC)
   LOOP %>
<LI>
Item = <A HREF="<%= item.catalog_url %>"><%= item.product_name %></A><BR>
Price = <BIG><%= item.list_price %></BIG><BR>
<% END LOOP; %>
</UL>
</BODY>
</HTML>
```

After loading Example 10–10 into the database, you can pass a parameter to the `show_product_partial` procedure through a URL. This example specifies a minimum price of 250:

```
http://www.example.com/pls/show_product_partial?p_minprice=250
```

Filtering results is appropriate for applications such as search results, where users might be overwhelmed by choices. But in a retail situation, you might want to use the alternative technique illustrated in Example 10–11, so that customers can still choose to purchase other items:

- Instead of filtering the results through a `WHERE` clause, retrieve the entire result set and then take different actions for different returned rows.
- Change the HTML to highlight the output that meets their criteria. Example 10–11 uses the background color for an HTML table row. You can also insert a special icon, increase the font size, or use another technique to call attention to the most important rows.
- Present the results in an HTML table.

Example 10–11 show_product_highlighed.psp

```
<%@ plsql procedure="show_product_highlighted" %>
<%@ plsql parameter="p_minprice" default="100" %>
<%! v_color VARCHAR2(7); %>

<HTML>
<HEAD><TITLE>Show Items Greater Than Specified Price</TITLE></HEAD>
<BODY>
<P>This report shows all items, highlighting those whose price is
 greater than <%= p_minprice %>.
<P>
<TABLE BORDER>
  <TR>
    <TH>Product</TH>
    <TH>Price</TH>
  </TR>
  <% FOR ITEM IN (SELECT product_name, list_price, catalog_url
                  FROM product_information
                  WHERE list_price IS NOT NULL
                  ORDER BY list_price DESC) LOOP
     IF item.list_price > p_minprice THEN
        v_color := '#CCCCFF';
     ELSE
        v_color := '#CCCCCC';
     END IF;
  %>
  <TR BGCOLOR="<%= v_color %>">
    <TD><A HREF="<%= item.catalog_url %>"><%= item.product_name %></A></TD>
    <TD><BIG><%= item.list_price %></BIG></TD>
  </TR>
  <% END LOOP; %>
</TABLE>
</BODY>
</HTML>
```

Using an HTML Form to Invoke a PL/SQL Server Page

Example 10–12 shows a bare-bones HTML form that allows the user to enter a price. The form invokes the `show_product_partial` stored procedure illustrated in Example 10–10 and passes it the entered value as the `p_minprice` parameter.

To avoid coding the entire URL of the stored procedure in the `ACTION=` attribute of the form, you can make the form a PSP file so that it resides in the same directory as the PSP file that it invokes. Even though this HTML file contains no PL/SQL code, you

can give it a .psp extension and load it as a stored procedure into the database. When the product_form stored procedure is executed through a URL, it displays the HTML exactly as it appears in the file.

Example 10–12 product_form.psp

```
<HTML>
<BODY>
<FORM method="POST" action="show_product_partial">
  <P>Enter the minimum price you want to pay:
  <INPUT type="text" name="p_minprice">
  <INPUT type="submit" value="Submit">
</FORM>
</BODY>
</HTML>
```

Including JavaScript in a PSP File

To produce an elaborate HTML file, perhaps including dynamic content such as JavaScript, you can simplify the source code by implementing it as a PSP. This technique avoids having to deal with nested quotation marks, escape characters, concatenated literals and variables, and indentation of the embedded content.

Example 10–13 shows a version of Example 10–10 that uses JavaScript to display the order status in the browser status bar when the user moves his or her mouse over the product URL.

Example 10–13 show_product_javascript.psp

```
<%@ plsql procedure="show_product_javascript" %>
<%@ plsql parameter="p_minprice" default="100" %>
<HTML>
<HEAD>
  <TITLE>Show Items Greater Than Specified Price</TITLE>

<SCRIPT language="JavaScript">
<!--hide

var text=" ";

function overlink (text)
{
  window.status=text;
}
function offlink (text)
{
  window.status=text;
}

//-->
</SCRIPT>

</HEAD>
<BODY>
<P>This report shows the items whose price is greater than <%= p_minprice %>.
<P>
<UL>
<% FOR ITEM IN (SELECT product_name, list_price, catalog_url, product_status
                FROM product_information
                WHERE list_price > p_minprice
```

```
                ORDER BY list_price DESC)
   LOOP %>
<LI>
Item =
  <A HREF="<%= item.catalog_url %>"
  onMouseover="overlink('PRODUCT STATUS: <%= item.product_status %>');return true"
  onMouseout="offlink(' ');return true">
    <%= item.product_name %>
  </A>
<BR>
Price = <BIG><%= item.list_price %></BIG><BR>
<% END LOOP; %>
</UL>
</BODY>
</HTML>
```

Debugging PL/SQL Server Pages

As you begin experimenting with PL/SQL Server Pages, and as you adapt your first simple pages into more elaborate ones, keep these guidelines in mind when you encounter problems:

- The first step is to get all the PL/SQL syntax and PSP directive syntax right. If you make a mistake here, the file does not compile.
  - Use semicolons to terminate lines where required.
  - If a value must be quoted, quote it. You might need to enclose a value in single quotation marks (which PL/SQL needs) inside double quotation marks (which PSP needs).
  - Mistakes in the PSP directives are usually reported through PL/SQL syntax messages. Check that your directives use the right syntax, that directives are closed properly, and that you are using the right element (declaration, expression, or code block) depending on what goes inside it.
  - PSP attribute names are case-sensitive. Most are specified in all lowercase; `contentType` and `errorPage` must be specified as mixed-case.
- When using a URL to request a PSP, you might get an error that the file is not found. In this case:
  - Be sure you are requesting the right virtual path, depending on the way the web gateway is configured. Typically, the path includes the host name, optionally a port number, the schema name, and the name of the stored procedure (with no `.psp` extension).
  - If you use the `-replace` option when compiling the file, the old version of the stored procedure is erased. So, after a failed compilation, you must fix the error or the page is not available. You might want to test scripts in a separate schema, then load them into the production schema.
  - If you copied the file from another file, remember to change any procedure name directives in the source to match the correct file name.
  - When you get one file-not-found error, request the latest version of the page the next time. The error page might be cached by the browser. You might need to force a page reload in the browser to bypass the cache.
- When the PSP script is run, and the results come back to the browser, use standard debugging techniques to check for and correct wrong output. The difficult part is to configure the interface between different HTML forms, scripts, and CGI

programs so that the right values are passed into your page. The page might return an error because of a parameter mismatch.

Guidelines:

- To determine exactly what is being passed to your page, use `METHOD=GET` in the invoking form so that the parameters are visible in the URL.
- Ensure that the form or CGI program that invokes your page passes the correct number of parameters, and that the names specified by the `NAME=` attributes on the form match the parameter names in the PSP file. If the form includes any hidden input fields, or uses the `NAME=` attribute on the `Submit` or `Reset` buttons, then the PSP file must declare equivalent parameters.
- Ensure that the parameters can be cast from string into the correct PL/SQL types. For example, do not include alphabetic characters if the parameter in the PSP file is declared as a `NUMBER`.
- Ensure that the query string of the URL consists of name-value pairs, separated by equals signs, especially if you are passing parameters by constructing a hard-coded link to the page.
- If you are passing a lot of parameter data, such as large strings, you might exceed the volume that can be passed with `METHOD=GET`. You can switch to `METHOD=POST` in the invoking form without changing your PSP file.
- Although the `loadpsp` command reports line numbers correctly when there is a syntax error in your source file, line numbers reported for run-time errors refer to a transformed version of the source and do not match the line numbers in the original source. When you encounter errors that produce an error trace instead of the expected web page, you must locate the error through exception handlers and by printing debug output.

Putting PL/SQL Server Pages into Production

Before putting your PSP application into production, consider issues such as usability and download speed:

- Pages can be rendered faster in the browser if the `HEIGHT=` and `WIDTH=` attributes are specified for all images. You might standardize on picture sizes, or store the height and width of images in the database along with the data or URL.
- For viewers who turn off graphics, or who use alternative browsers that read the text out loud, include a description of significant images using the `ALT=` attribute. You might store the description in the database along with the image.
- Although an HTML table provides a good way to display data, a large table can make your application seem slow. Often, the reader sees a blank page until the entire table is downloaded. If the amount of data in an HTML table is large, consider splitting the output into multiple tables.
- If you set text, font, or background colors, test your application with different combinations of browser color settings:
  - Test what happens if you override just the foreground color in the browser, or just the background color, or both.
  - If you set one color (such as the foreground text color), set all the colors through the `<BODY>` tag, to avoid hard-to-read combinations like white text on a white background.

  - If you use a background image, specify a similar background color to provide proper contrast for viewers who do not load graphics.
  - If the information conveyed by different colors is crucial, consider using an alternative technique. For example, you might put an icon next to special items in a table. Some users might see your page on a monochrome screen or on browsers that cannot represent different colors.
- Providing context information prevents users from getting lost. Include a descriptive `<TITLE>` tag for your page. If the user is partway through a procedure, indicate which step is represented by your page. Provide links to logical points to continue with the procedure, return to a previous step, or cancel the procedure completely. Many pages might use a standard set of links that you embed using the include directive.
- In any entry fields, users might enter incorrect values. Where possible, use `SELECT` lists to present a set of choices. Validate any text entered in a field before passing it to SQL. The earlier you can validate, the better; a JavaScript function can detect incorrect data and prompt the user to correct it before they press the `Submit` button and call the database.
- Browsers tend to be lenient when displaying incorrect HTML. What looks OK in one browser might look bad or might not display at all in another browser.

  Guidelines:

  - Pay attention to HTML rules for quotation marks, closing tags, and especially for anything to do with tables.
  - Minimize the dependence on tags that are only supported by a single browser. Sometimes you can provide an extra bonus using such tags, but your application must still be usable with other browsers.
  - You can check the validity, and even in some cases the usability, of your HTML for free at many sites on the World Wide Web.

11

Using Continuous Query Notification (CQN)

Continuous Query Notification (CQN) allows an application to register queries with the database for either object change notification (the default) or query result change notification. An object referenced by a registered query is a **registered object**.

If a query is registered for **object change notification (OCN)**, the database notifies the application whenever a transaction changes an object that the query references and commits, regardless of whether the query result changed.

If a query is registered for **query result change notification (QRCN)**, the database notifies the application whenever a transaction changes the result of the query and commits.

A **CQN registration** associates a list of one or more queries with a notification type (OCN or QRCN) and a notification handler. To create a CQN registration, you can use either the PL/SQL interface or Oracle Call Interface (OCI). If you use the PL/SQL interface, the notification handler is a server-side PL/SQL stored procedure; if you use OCI, the notification handler is a client-side C callback procedure.

This chapter explains general CQN concepts and explains how to use the PL/SQL CQN interface. For information about using OCI for CQN, see *Oracle Call Interface Programmer's Guide*.

Topics:

- Object Change Notification (OCN)
- Query Result Change Notification (QRCN)
- Events that Generate Notifications
- Notification Contents
- Good Candidates for CQN
- Creating CQN Registrations
- Querying CQN Registrations
- Interpreting Notifications
- Deleting Registrations
- Configuring CQN: Scenario

Note: The terms **OCN** and **QRCN** refer to both the notification type and the notification itself: An application registers a query *for OCN*, and the database sends the application *an OCN*; an application registers a query *for QRCN*, and the database sends the application *a QRCN*.

Object Change Notification (OCN)

If an application registers a query for object change notification (OCN), the database sends the application an OCN whenever a transaction changes an object associated with the query and commits, regardless of whether the result of the query changed.

For example, if an application registers the query in Example 11–1 for OCN, and a user commits a transaction that changes the `EMPLOYEES` table, the database sends the application an OCN, even if the changed row or rows did not satisfy the query predicate (for example, if `DEPARTMENT_ID` = 5).

Example 11–1 Query to be Registered for Change Notification

```
SELECT EMPLOYEE_ID, SALARY
FROM EMPLOYEES
WHERE DEPARTMENT_ID = 10;
```

Query Result Change Notification (QRCN)

Note: For QRCN support, the `COMPATIBLE` initialization parameter of the database must be at least 11.0.0, and Automatic Undo Management (AUM) must be enabled (as it is by default).

For information about the `COMPATIBLE` initialization parameter, see *Oracle Database Administrator's Guide*.

For information about AUM, see *Oracle Database Administrator's Guide*.

If an application registers a query for query result change notification (QRCN), the database sends the application a QRCN whenever a transaction changes the result of the query and commits.

For example, if an application registers the query in Example 11–1 for QRCN, the database sends the application a QRCN only if the query result set changes; that is, if one of these data manipulation language (DML) statements commits:

- An `INSERT` or `DELETE` of a row that satisfies the query predicate (`DEPARTMENT_ID` = 10).
- An `UPDATE` to the `EMPLOYEE_ID` or `SALARY` column of a row that satisfied the query predicate (`DEPARTMENT_ID` = 10).
- An `UPDATE` to the `DEPARTMENT_ID` column of a row that changed its value from 10 to a value other than 10, causing the row to be deleted from the result set.
- An `UPDATE` to the `DEPARTMENT_ID` column of a row that changed its value to 10 from a value other than 10, causing the row to be added to the result set.

The default notification type is OCN. For QRCN, specify `QOS_QUERY` in the `QOSFLAGS` attribute of the `CQ_NOTIFICATION$_REG_INFO` object.

With QRCN, you have a choice of guaranteed mode (the default) or best-effort mode.

Topics:

- Guaranteed Mode
- Best-Effort Mode

Guaranteed Mode

In guaranteed mode, there are no false positives: the database sends the application a QRCN only when the query result set is guaranteed to have changed.

For example, suppose that an application registered the query in Example 11–1 for QRCN, that employee 201 is in department 10, and that these statements are executed:

```
UPDATE EMPLOYEES
SET SALARY = SALARY + 10
WHERE EMPLOYEE_ID = 201;

UPDATE EMPLOYEES
SET SALARY = SALARY - 10
WHERE EMPLOYEE_ID = 201;

COMMIT;
```

Each `UPDATE` statement in the preceding transaction changes the query result set, but together they have no effect on the query result set; therefore, the database does not send the application a QRCN for the transaction.

For guaranteed mode, specify `QOS_QUERY`, but not `QOS_BEST_EFFORT`, in the `QOSFLAGS` attribute of the `CQ_NOTIFICATION$_REG_INFO` object.

Some queries are too complex for QRCN in guaranteed mode. For the characteristics of queries that can be registered in guaranteed mode, see "Queries that Can Be Registered for QRCN in Guaranteed Mode" on page 11-15.

Best-Effort Mode

Some queries that are too complex for guaranteed mode can be registered for QRCN in best-effort mode, in which CQN creates and registers simpler versions of them.

For example, the query in Example 11–2 is too complex for QRCN in guaranteed mode because it contains the aggregate function `SUM`.

Example 11–2 Query Too Complex for QRCN in Guaranteed Mode

```
SELECT SUM(SALARY)
FROM EMPLOYEES
WHERE DEPARTMENT_ID = 20;
```

In best-effort mode, CQN registers this simpler version of the query in Example 11–2:

```
SELECT SALARY
FROM EMPLOYEES
WHERE DEPARTMENT_ID = 20;
```

Whenever the result of the original query changes, the result of its simpler version also changes; therefore, no notifications are lost from the simplification. However, the simplification might cause false positives, because the result of the simpler version can change when the result of the original query does not.

In best-effort mode, the database:

- Minimizes the OLTP response overhead that is from notification-related processing, as follows:
  - For a single-table query, the database determines whether the query result has changed by which columns changed and which predicates the changed rows satisfied.
  - For a multiple-table query (a join), the database uses the primary-key/foreign-key constraint relationships between the tables to determine whether the query result has changed.
- Sends the application a QRCN whenever a DML statement changes the query result set, even if a subsequent DML statement nullifies the change made by the first DML statement.

The overhead minimization of best-effort mode infrequently causes false positives, even for queries that CQN does not simplify. For example, consider the query in Example 11–1 and the transaction in "Guaranteed Mode" on page 11-3. In best-effort mode, CQN does not simplify the query, but the transaction generates a false positive.

Some types of queries are so simplified that invalidations are generated at object level; that is, whenever any object referenced in those queries changes. Examples of such queries are those that use unsupported column types or include subqueries. The solution to this problem is to rewrite the original queries.

For example, the query in Example 11–3 is too complex for QRCN in guaranteed mode because it includes a subquery.

Example 11–3 Query Whose Simplified Version Invalidates Objects

```
SELECT SALARY
FROM EMPLOYEES
WHERE DEPARTMENT_ID IN (
  SELECT DEPARTMENT_ID
  FROM DEPARTMENTS
  WHERE LOCATION_ID = 1700
);
```

In best-effort mode, CQN simplifies the query in Example 11–3 to this:

```
SELECT * FROM EMPLOYEES, DEPARTMENTS;
```

The simplified query can cause objects to be invalidated. However, if you rewrite the original query as follows, you can register it in either guaranteed mode or best-effort mode:

```
SELECT SALARY
FROM EMPLOYEES, DEPARTMENTS
WHERE EMPLOYEES.DEPARTMENT_ID = DEPARTMENTS.DEPARTMENT_ID
  AND DEPARTMENTS.LOCATION_ID = 1700;
```

Queries that can be registered only in best-effort mode are described in "Queries that Can Be Registered for QRCN Only in Best-Effort Mode" on page 11-16.

The default for QRCN mode is guaranteed mode. For best-effort mode, specify `QOS_BEST_EFFORT` in the `QOSFLAGS` attribute of the `CQ_NOTIFICATION$_REG_INFO` object.

Events that Generate Notifications

These events generate notifications:

- Committed DML Transactions
- Committed DDL Statements
- Deregistration
- Global Events

Committed DML Transactions

When the notification type is OCN, any DML transaction that changes one or more registered objects generates one notification for each object when it commits.

When the notification type is QRCN, any DML transaction that changes the result of one or more registered queries generates a notification when it commits. The notification includes the query IDs of the queries whose results changed.

For either notification type, the notification includes:

- Name of each changed table
- Operation type (`INSERT`, `UPDATE`, or `DELETE`)
- `ROWID` of each changed row, if the registration was created with the `ROWID` option and the number of modified rows was not too large. For more information, see "ROWID Option" on page 11-12.

Committed DDL Statements

For both OCN and QRCN, these data definition language (DDL) statements, when committed, generate notifications:

- `ALTER TABLE`
- `TRUNCATE TABLE`
- `FLASHBACK TABLE`
- `DROP TABLE`

Note: When the notification type is OCN, a committed `DROP TABLE` statement generates a `DROP NOTIFICATION`.

Any OCN registrations of queries on the dropped table become disassociated from that table (which no longer exists), but the registrations themselves continue to exist. If any of these registrations are associated with objects other than the dropped table, committed changes to those other objects continue to generate notifications. Registrations associated only with the dropped table also continue to exist, and their creator can add queries (and their referenced objects) to them.

An OCN registration is based on the version and definition of an object at the time the query was registered. If an object is dropped, registrations on that object are disassociated from it forever. If an object is created with the same name, and in the same schema, as the dropped object, the created object is not associated with OCN registrations that were associated with the dropped object.

When the notification type is QRCN:

- The notification includes:
  - Query IDs of the queries whose results have changed
  - Name of the modified table
  - Type of DDL operation
- Some DDL operations that invalidate registered queries can cause those queries to be deregistered.

  For example, suppose that this query is registered for QRCN:

  ```
  SELECT COL1 FROM TEST_TABLE
    WHERE COL2 = 1;
  ```

  Suppose that `TEST_TABLE` has this schema:

  ```
  (COL1 NUMBER, COL2 NUMBER, COL3 NUMBER)
  ```

  This DDL statement, when committed, invalidates the query and causes it to be removed from the registration:

  ```
  ALTER TABLE DROP COLUMN COL2;
  ```

Deregistration

For both OCN and QRCN, deregistration—removal of a registration from the database—generates a notification. The reasons that the database removes a registration are:

- Timeout

  If `TIMEOUT` is specified with a nonzero value when the queries are registered, the database purges the registration after the specified time interval.

  If `QOS_DEREG_NFY` is specified when the queries are registered, the database purges the registration after it generates its first notification.
- Loss of privileges

  If privileges are lost on an object associated with a registered query, and the notification type is OCN, the database purges the registration. (When the notification type is QRCN, the database removes that query from the registration, but does not purge the registration.)

  For privileges needed to register queries, see "Prerequisites for Creating CQN Registrations" on page 11-14.

A notification is not generated when a client application performs an explicit deregistration.

Global Events

The global events `EVENT_STARTUP` and `EVENT_SHUTDOWN` generate notifications.

In an Oracle RAC environment, these events generate notifications:

- `EVENT_STARTUP` when the first instance of the database starts
- `EVENT_SHUTDOWN` when the last instance of the database shuts down
- `EVENT_SHUTDOWN_ANY` when any instance of the database shuts down

The preceding global events are constants defined in the `DBMS_CQ_NOTIFICATION` package.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for more information about the `DBMS_CQ_NOTIFICATION` package

Notification Contents

A notification contains some or all of this information:

- Type of event, which is one of:
  - Startup
  - Object change
  - Query result change
  - Deregistration
  - Shutdown
- Registration ID of affected registration
- Names of changed objects
- If `ROWID` option was specified, `ROWID`s of changed rows
- If the notification type is QRCN: Query IDs of queries whose results changed
- If notification resulted from a DML or DDL statement:
  - Array of names of modified tables
  - Operation type (for example, `INSERT` or `UPDATE`)

A notification does not contain the changed data itself. For example, the notification does not say that a monthly salary increased from 5000 to 6000. To obtain more recent values for the changed objects or rows or query results, the application must query the database.

Good Candidates for CQN

Good candidates for CQN are applications that cache the result sets of queries on infrequently changed objects in the middle tier, to avoid network round trips to the database. These applications can use CQN to register the queries to be cached. When such an application receives a notification, it can refresh its cache by rerunning the registered queries.

An example of such an application is a web forum. Because its users need not view content as soon as it is inserted into the database, this application can cache information in the middle tier and have CQN tell it when it when to refresh the cache.

Figure 11–1 illustrates a typical scenario in which the database serves data that is cached in the middle tier and then accessed over the Internet.

Figure 11–1 Middle-Tier Caching

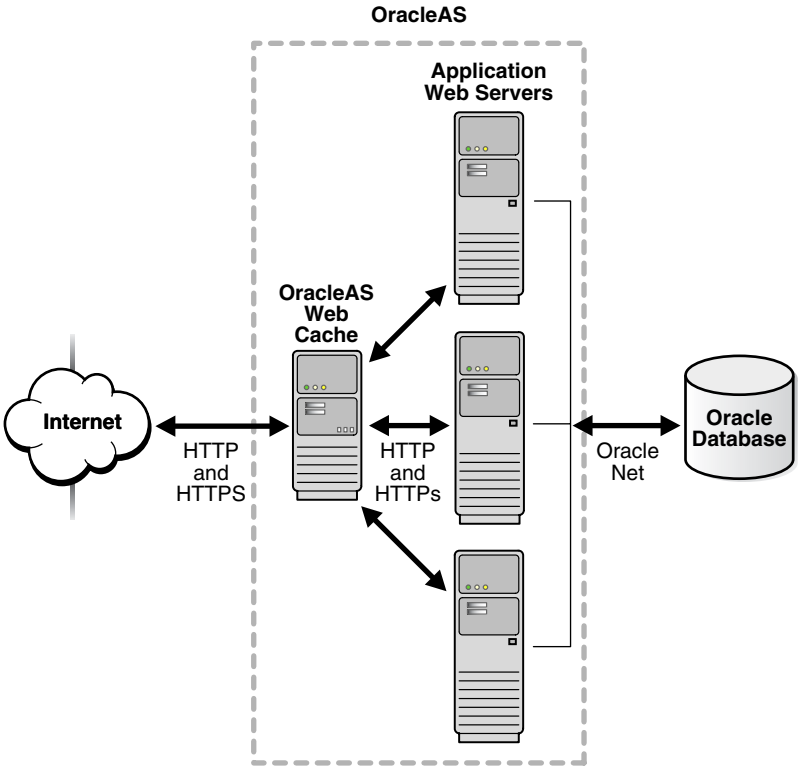

Applications in the middle tier require rapid access to cached copies of database objects while keeping the cache as current as possible in relation to the database. Cached data becomes obsolete when a transaction modifies the data and commits, thereby putting the application at risk of accessing incorrect results. If the application uses CQN, the database can publish a notification when a change occurs to registered objects with details on what changed. In response to the notification, the application can refresh cached data by fetching it from the back-end database.

Figure 11–2 illustrates the process by which middle-tier Web clients receive and process notifications.

Figure 11–2 Basic Process of Continuous Query Notification (CQN)

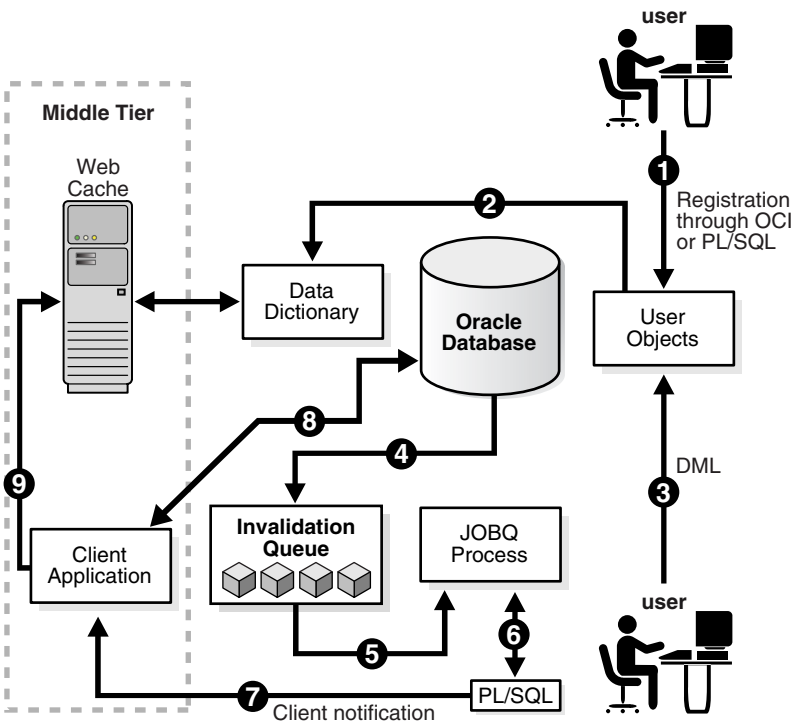

Explanation of steps in Figure 11–2 (if registrations are created using PL/SQL and that the application has cached the result set of a query on `HR.EMPLOYEES`):

1. The developer uses PL/SQL to create a CQN registration for the query, which consists of creating a stored PL/SQL procedure to process notifications and then using the PL/SQL CQN interface to create a registration for the query, specifying the PL/SQL procedure as the notification handler.
2. The database populates the registration information in the data dictionary.
3. A user updates a row in the `HR.EMPLOYEES` table in the back-end database and commits the update, causing the query result to change. The data for `HR.EMPLOYEES` cached in the middle tier is now outdated.
4. The database adds a message that describes the change to an internal queue.
5. The database notifies a `JOBQ` background process of a notification message.
6. The `JOBQ` process runs the stored procedure specified by the client application. In this example, `JOBQ` passes the data to a server-side PL/SQL procedure. The implementation of the PL/SQL notification handler determines how the notification is handled.
7. Inside the server-side PL/SQL procedure, the developer can implement logic to notify the middle-tier client application of the changes to the registered objects. For example, it notifies the application of the `ROWID` of the changed row in `HR.EMPLOYEES`.
8. The client application in the middle tier queries the back-end database to retrieve the data in the changed row.
9. The client application updates the cache with the data.

Creating CQN Registrations

A **CQN registration** associates a list of one or more queries with a notification type and a notification handler.

The notification type is either OCN or QRCN. For information about these types, see "Object Change Notification (OCN)" on page 11-2 and "Query Result Change Notification (QRCN)" on page 11-2.

To create a CQN registration, you can use either the PL/SQL interface or OCI. If you use the PL/SQL interface, the notification handler is a server-side PL/SQL stored procedure; if you use OCI, the notification handler is a client-side C callback procedure. (This topic explains only the PL/SQL interface. For information about OCI, see *Oracle Call Interface Programmer's Guide*.)

Once created, a registration is stored in the database. In an Oracle RAC environment, it is visible to all database instances. Transactions that change the query results in any database instance generate notifications.

By default, a registration survives until the application that created it explicitly deregisters it or until the database implicitly purges it (from loss of privileges, for example).

Topics:

- PL/SQL CQN Registration Interface
- CQN Registration Options
- Prerequisites for Creating CQN Registrations
- Queries that Can Be Registered for Object Change Notification (OCN)
- Queries that Can Be Registered for Query Result Change Notification (QRCN)
- Using PL/SQL to Register Queries for CQN
- Best Practices for CQN Registrations
- Troubleshooting CQN Registrations

PL/SQL CQN Registration Interface

The PL/SQL CQN registration interface is implemented with the `DBMS_CQ_NOTIFICATION` package. You use the `DBMS_CQ_NOTIFICATION.NEW_REG_START` function to open a registration block. You specify the registration details, including the notification type and notification handler, as part of the `CQ_NOTIFICATION$_REG_INFO` object, which is passed as an argument to the `NEW_REG_START` procedure. Every query that you run while the registration block is open is registered with CQN. If you specified notification type QRCN, the database assigns a query ID to each query. You can retrieve these query IDs with the `DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID` function. To close the registration block, you use the `DBMS_CQ_NOTIFICATION.REG_END` function.

For step-by-step instructions, see "Using PL/SQL to Register Queries for CQN" on page 11-18.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for more information about the `DBMS_CQ_NOTIFICATION` package

CQN Registration Options

You can change the CQN registration defaults with the options summarized in Table 11–1.

Table 11–1 Continuous Query Notification Registration Options

Option	Description
Notification Type	Specifies QRCN (the default is OCN).
QRCN Mode[1]	Specifies best-effort mode (the default is guaranteed mode).
`ROWID`	Includes the value of the `ROWID` pseudocolumn for each changed row in the notification.
Operations Filter[2]	Publishes the notification only if the operation type matches the specified filter condition.
Transaction Lag[2]	Deprecated. Use Notification Grouping instead.
Notification Grouping	Specifies how notifications are grouped.
Reliable	Stores notifications in a persistent database queue (instead of in shared memory, the default).
Purge on Notify	Purges the registration after the first notification.
Timeout	Purges the registration after a specified time interval.

[1] Applies only when notification type is QRCN.
[2] Applies only when notification type is OCN.

Topics:

- Notification Type Option
- QRCN Mode (QRCN Notification Type Only)
- ROWID Option
- Operations Filter Option (OCN Notification Type Only)
- Transaction Lag Option (OCN Notification Type Only)
- Notification Grouping Options
- Reliable Option
- Purge-on-Notify and Timeout Options

Notification Type Option

The notification types are OCN (described in "Object Change Notification (OCN)" on page 11-2) and QRCN (described in "Query Result Change Notification (QRCN)" on page 11-2).

QRCN Mode (QRCN Notification Type Only)

The QRCN mode option applies only when the notification type is QRCN. Instructions for setting the notification type to QRCN are in "Notification Type Option" on page 11-11.

The QRCN modes are guaranteed (described in "Guaranteed Mode" on page 11-3) and best-effort (described in "Best-Effort Mode" on page 11-3).

The default is guaranteed mode. For best-effort mode, specify `QOS_BEST_EFFORT` in the `QOSFLAGS` attribute of the `CQ_NOTIFICATION$_REG_INFO` object.

ROWID Option

The ROWID option includes the value of the ROWID pseudocolumn (the rowid of the row) for each changed row in the notification. To include the ROWID option of each changed row in the notification, specify QOS_ROWIDS in the QOSFLAGS attribute of the CQ_NOTIFICATION$_REG_INFO object.

From the ROWID information in the notification, the application can retrieve the contents of the changed rows by performing queries of this form:

```
SELECT * FROM table_name_from_notification
WHERE ROWID = rowid_from_notification;
```

ROWIDs are published in the external string format. For a regular heap table, the length of a ROWID is 18 character bytes. For an Index Organized Table (IOT), the length of the ROWID depends on the size of the primary key, and might exceed 18 bytes.

If the server does not have enough memory for the ROWIDs, the notification might be "rolled up" into a FULL-TABLE-NOTIFICATION, indicated by a special flag in the notification descriptor. Possible reasons for a FULL-TABLE-NOTIFICATION are:

- Total shared memory consumption from ROWIDs exceeds 1% of the dynamic shared pool size.
- Too many rows were changed in a single registered object within a transaction (the upper limit is approximately 80).
- Total length of the logical ROWIDs of modified rows for an IOT is too large (the upper limit is approximately 1800 bytes).
- You specified the Notification Grouping option NTFN_GROUPING_TYPE with the value DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_SUMMARY, described in "Notification Grouping Options" on page 11-13.

Because a FULL-TABLE-NOTIFICATION does not include ROWIDs, the application that receives it must assume that the entire table (that is, all rows) might have changed.

Operations Filter Option (OCN Notification Type Only)

The Operations Filter option applies only when the notification type is OCN.

The Operations Filter option enables you to specify the types of operations that generate notifications.

The default is all operations. To specify that only some operations generate notifications, use the OPERATIONS_FILTER attribute of the CQ_NOTIFICATION$_REG_INFO object. With the OPERATIONS_FILTER attribute, specify the type of operation with the constant that represents it, which is defined in the DBMS_CQ_NOTIFICATIONS package, as follows:

Operation	Constant
INSERT	DBMS_CQ_NOTIFICATIONS.INSERTOP
UPDATE	DBMS_CQ_NOTIFICATIONS.UPDATEOP
DELETE	DBMS_CQ_NOTIFICATIONS.DELETEOP
ALTEROP	DBMS_CQ_NOTIFICATIONS.ALTEROP
DROPOP	DBMS_CQ_NOTIFICATIONS.DROPOP
UNKNOWNOP	DBMS_CQ_NOTIFICATIONS.UNKNOWNOP
All (default)	DBMS_CQ_NOTIFICATIONS.ALL_OPERATIONS

To specify multiple operations, use bitwise OR. For example:

```
DBMS_CQ_NOTIFICATIONS.INSERTOP + DBMS_CQ_NOTIFICATIONS.DELETEOP
```

OPERATIONS_FILTER has no effect if you also specify QOS_QUERY in the QOSFLAGS attribute, because QOS_QUERY specifies notification type QRCN.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for more information about the DBMS_CQ_NOTIFICATION package

Transaction Lag Option (OCN Notification Type Only)

The Transaction Lag option applies only when the notification type is OCN.

> **Note:** This option is deprecated. To implement flow-of-control notifications, use "Notification Grouping Options" on page 11-13.

The Transaction Lag option specifies the number of transactions by which the client application can lag behind the database. If the number is 0, every transaction that changes a registered object results in a notification. If the number is 5, every fifth transaction that changes a registered object results in a notification. The database tracks intervening changes at object granularity and includes them in the notification, so that the client does not lose them.

A transaction lag greater than 0 is useful only if an application implements flow-of-control notifications. Ensure that the application generates notifications frequently enough to satisfy the lag, so that they are not deferred indefinitely.

If you specify TRANSACTION_LAG, then notifications do not include ROWIDs, even if you also specified QOS_ROWIDS.

Notification Grouping Options

By default, notifications are generated immediately after the event that causes them.

Notification Grouping options, which are attributes of the CQ_NOTIFICATION$_REG_INFO object, are:

Attribute	Description
NTFN_GROUPING_CLASS	Specifies the class by which to group notifications. The only allowed values are DBMS_CQ_NOTIFICATION.NTFN_GROUPING_CLASS_TIME, which groups notifications by time, and zero, which is the default (notifications are generated immediately after the event that causes them).
NTFN_GROUPING_VALUE	Specifies the time interval that defines the group, in seconds. For example, if this value is 900, notifications generated in the same 15-minute interval are grouped.
NTFN_GROUPING_TYPE	Specifies the type of grouping, which is either of: ▪ DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_SUMMARY: All notifications in the group are summarized into a single notification. **Note:** The single notification does not include ROWIDs, even if you specified the ROWID option. ▪ DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_LAST: Only the last notification in the group is published and the earlier ones discarded.

Attribute	Description
NTFN_GROUPING_START_TIME	Specifies when to start generating notifications. If specified as NULL, it defaults to the current system-generated time.
NTFN_GROUPING_REPEAT_COUNT	Specifies how many times to repeat the notification. Set to DBMS_CQ_NOTIFICATION.NTFN_GROUPING_FOREVER to receive notifications for the life of the registration. To receive at most *n* notifications during the life of the registration, set to *n*.

Note: Notifications generated by timeouts, loss of privileges, and global events might be published before the specified grouping interval expires. If they are, any pending grouped notifications are also published before the interval expires.

Reliable Option

By default, a CQN registration is stored in shared memory. To store it in a persistent database queue instead—that is, to generate **reliable notifications**—specify QOS_RELIABLE in the QOSFLAGS attribute of the CQ_NOTIFICATION$_REG_INFO object.

The advantage of reliable notifications is that if the database fails after generating them, it can still deliver them after it restarts. In an Oracle RAC environment, a surviving database instance can deliver them.

The disadvantage of reliable notifications is that they have higher CPU and I/O costs than default notifications do.

Purge-on-Notify and Timeout Options

By default, a CQN registration survives until the application that created it explicitly deregisters it or until the database implicitly purges it (from loss of privileges, for example).

To purge the registration after it generates its first notification, specify QOS_DEREG_NFY in the QOSFLAGS attribute of the CQ_NOTIFICATION$_REG_INFO object.

To purge the registration after *n* seconds, specify *n* in the TIMEOUT attribute of the CQ_NOTIFICATION$_REG_INFO object.

You can use the Purge-on-Notify and Timeout options together.

Prerequisites for Creating CQN Registrations

These are prerequisites for creating CQN registrations:

- You must have these privileges:
  - EXECUTE privilege on the DBMS_CQ_NOTIFICATION package, whose subprograms you use to create a registration
  - CHANGE NOTIFICATION system privilege
  - SELECT privileges on all objects to be registered

  Loss of privileges on an object associated with a registered query generates a notification—see "Deregistration" on page 11-6.
- You must be connected as a non-SYS user.

- You must not be in the middle of an uncommitted transaction.
- The `dml_locks` `init.ora` parameter must have a nonzero value (as its default value does).

  (This is also a prerequisite for receiving notifications.)

> **Note:** For QRCN support, the `COMPATIBLE` setting of the database must be at least 11.0.0.

Queries that Can Be Registered for Object Change Notification (OCN)

Most queries can be registered for OCN, including those executed as part of stored procedures and `REF` cursors.

Queries that cannot be registered for OCN are:

- Queries on fixed tables or fixed views
- Queries on user views
- Queries that contain database links (dblinks)
- Queries over materialized views

> **Note:** You can use synonyms in OCN registrations, but not in QRCN registrations.

Queries that Can Be Registered for Query Result Change Notification (QRCN)

Some queries can be registered for QRCN in guaranteed mode, some can be registered for QRCN only in best-effort mode, and some cannot be registered for QRCN in either mode. (For information about modes, see "Guaranteed Mode" on page 11-3 and "Best-Effort Mode" on page 11-3.)

Topics:

- Queries that Can Be Registered for QRCN in Guaranteed Mode
- Queries that Can Be Registered for QRCN Only in Best-Effort Mode
- Queries that Cannot Be Registered for QRCN in Either Mode

Queries that Can Be Registered for QRCN in Guaranteed Mode

To be registered for QRCN in guaranteed mode, a query must conform to these rules:

- Every column that it references is either a `NUMBER` data type or a `VARCHAR2` data type.
- Arithmetic operators in column expressions are limited to these binary operators, and their operands are columns with numeric data types:
  - `+` (addition)
  - `-` (subtraction, not unary minus)
  - `*` (multiplication)
  - `/` (division)
- Comparison operators in the predicate are limited to:

  - < (less than)
  - <= (less than or equal to)
  - = (equal to)
  - >= (greater than or equal to)
  - > (greater than)
  - <> or != (not equal to)
  - IS NULL
  - IS NOT NULL
- Boolean operators in the predicate are limited to AND, OR, and NOT.
- The query contains no aggregate functions (such as SUM, COUNT, AVERAGE, MIN, and MAX).

  For a list of built-in SQL aggregate functions, see *Oracle Database SQL Language Reference*.

Guaranteed mode supports most queries on single tables and some inner equijoins, such as:

```
SELECT SALARY FROM EMPLOYEES, DEPARTMENTS
  WHERE EMPLOYEES.DEPARTMENT_ID = DEPARTMENTS.DEPARTMENT_ID
    AND DEPARTMENTS.LOCATION_ID = 1700;
```

Notes:

- Sometimes the query optimizer uses an execution plan that makes a query incompatible for guaranteed mode (for example, OR-expansion). For information about the query optimizer, see *Oracle Database Performance Tuning Guide*.
- Queries that can be registered in guaranteed mode can also be registered in best-effort mode, but results might differ, because best-effort mode can cause false positives even for queries that CQN does not simplify. For details, see "Best-Effort Mode" on page 11-3.

Queries that Can Be Registered for QRCN Only in Best-Effort Mode

A query that does any of the following can be registered for QRCN only in best-effort mode, and its simplified version generates notifications at object granularity:

- Refers to columns that have encryption enabled
- Has more than 10 items of the same type in the SELECT list
- Has expressions that include any of these:
  - String functions (such as SUBSTR, LTRIM, and RTRIM)
  - Arithmetic functions (such as TRUNC, ABS, and SQRT)

    For a list of built-in SQL functions, see *Oracle Database SQL Language Reference*.
  - Pattern-matching conditions LIKE and REGEXP_LIKE
  - EXISTS or NOT EXISTS condition

- Has disjunctions involving predicates defined on columns from different tables. For example:

```
SELECT EMPLOYEE_ID, DEPARTMENT_ID
  FROM EMPLOYEES, DEPARTMENTS
    WHERE EMPLOYEES.EMPLOYEE_ID = 10
      OR DEPARTMENTS.DEPARTMENT_ID = 'IT';
```

- Has user rowid access. For example:

```
SELECT DEPARTMENT_ID
  FROM DEPARTMENTS
    WHERE ROWID = 'AAANkdAABAAALinAAF';
```

- Has any join other than an inner join
- Has an execution plan that involves any of these:
  - Bitmap join, domain, or function-based indexes
  - `UNION ALL` or `CONCATENATION`

    (Either in the query itself, or as the result of an `OR`-expansion execution plan chosen by the query optimizer.)
  - `ORDER BY` or `GROUP BY`

    (Either in the query itself, or as the result of a `SORT` operation with an `ORDER BY` option in the execution plan chosen by the query optimizer.)
  - Partitioned index-organized table (IOT) with overflow segment
  - Clustered objects
  - Parallel execution

Queries that Cannot Be Registered for QRCN in Either Mode

A query that refers to any of the following cannot be registered for QRCN in either guaranteed or best-effort mode:

- Views
- Tables that are fixed, remote, or have Virtual Private Database (VPD) policies enabled
- `DUAL` (in the `SELECT` list)
- Synonyms
- Calls to user-defined PL/SQL subprograms
- Operators not listed in "Queries that Can Be Registered for QRCN in Guaranteed Mode" on page 11-15
- The aggregate function `COUNT`

  (Other aggregate functions are allowed in best-effort mode, but not in guaranteed mode.)
- Application contexts; for example:

```
SELECT SALARY FROM EMPLOYEES
WHERE USER = SYS_CONTEXT('USERENV', 'SESSION_USER');
```

- `SYSDATE`, `SYSTIMESTAMP`, or `CURRENT TIMESTAMP`

Also, a query that the query optimizer has rewritten using a materialized view cannot be registered for QRCN. For information about the query optimizer, see *Oracle Database Performance Tuning Guide*.

Using PL/SQL to Register Queries for CQN

To use PL/SQL to create a CQN registration, follow these steps:

1. Create a stored PL/SQL procedure to serve as the notification handler.
2. Create a `CQ_NOTIFICATION$_REG_INFO` object that specifies the name of the notification handler, the notification type, and other attributes of the registration.
3. In your client application, use the `DBMS_CQ_NOTIFICATION.NEW_REG_START` function to open a registration block.
4. Run the queries to register. (Do not run DML or DDL operations.)
5. Close the registration block, using the `DBMS_CQ_NOTIFICATION.REG_END` function.

Topics:

- Creating a PL/SQL Notification Handler
- Creating a CQ_NOTIFICATION$_REG_INFO Object
- Identifying Individual Queries in a Notification
- Adding Queries to an Existing Registration

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for more information about the `CQ_NOTIFICATION$_REG_INFO` object and the functions `NEW_REG_START` and `REG_END`, all of which are defined in the `DBMS_CQ_NOTIFICATION` package

Creating a PL/SQL Notification Handler

The PL/SQL stored procedure that you create to serve as the notification handler must have this signature:

```
PROCEDURE schema_name.proc_name(ntfnds IN CQ_NOTIFICATION$_DESCRIPTOR)
```

In the preceding signature, *`schema_name`* is the name of the database schema, *`proc_name`* is the name of the stored procedure, and *`ntfnds`* is the notification descriptor.

The notification descriptor is a `CQ_NOTIFICATION$_DESCRIPTOR` object, whose attributes describe the details of the change (transaction ID, type of change, queries affected, tables modified, and so on).

The `JOBQ` process passes the notification descriptor, *`ntfnds`*, to the notification handler, *`proc_name`*, which handles the notification according to its application requirements. (This is step 6 in Figure 11–2.)

> **Note:** The notification handler runs inside a job queue process. The `JOB_QUEUE_PROCESSES` initialization parameter specifies the maximum number of processes that can be created for the execution of jobs. You must set `JOB_QUEUE_PROCESSES` to a nonzero value to receive PL/SQL notifications.

Creating a CQ_NOTIFICATION$_REG_INFO Object

An object of type CQ_NOTIFICATION$_REG_INFO specifies the notification handler that the database runs when a registered objects changes. In SQL*Plus, you can view its type attributes by running this statement:

```
DESC CQ_NOTIFICATION$_REG_INFO
```

Table 11–2 describes the attributes of SYS.CQ_NOTIFICATION$_REG_INFO.

Table 11–2 Attributes of CQ_NOTIFICATION$_REG_INFO

Attribute	Description
CALLBACK	Specifies the name of the PL/SQL procedure to be executed when a notification is generated (a notification handler). You must specify the name in the form *schema_name.procedure_name*, for example, hr.dcn_callback.
QOSFLAGS	Specifies one or more quality-of-service flags, which are constants in the DBMS_CQ_NOTIFICATION package. For their names and descriptions, see Table 11–3. To specify multiple quality-of-service flags, use bitwise OR. For example: DBMS_CQ_NOTIFICATION.QOS_RELIABLE + DBMS_CQ_NOTIFICATION.QOS_ROWIDS
TIMEOUT	Specifies the timeout period for registrations. If set to a nonzero value, it specifies the time in seconds after which the database purges the registration. If 0 or NULL, then the registration persists until the client explicitly deregisters it. Can be combined with the QOSFLAGS attribute with its QOS_DEREG_NFY flag.
OPERATIONS_FILTER	Applies only to OCN (described in "Object Change Notification (OCN)" on page 11-2). Has no effect if you specify the QOS_FLAGS attribute with its QOS_QUERY flag. Filters messages based on types of SQL statement. You can specify these constants in the DBMS_CQ_NOTIFICATION package: ■ ALL_OPERATIONS notifies on all changes ■ INSERTOP notifies on inserts ■ UPDATEOP notifies on updates ■ DELETEOP notifies on deletes ■ ALTEROP notifies on ALTER TABLE operations ■ DROPOP notifies on DROP TABLE operations ■ UNKNOWNOP notifies on unknown operations You can specify a combination of operations with a bitwise OR. For example: DBMS_CQ_NOTIFICATION.INSERTOP + DBMS_CQ_NOTIFICATION.DELETEOP.

Table 11–2 (Cont.) Attributes of CQ_NOTIFICATION$_REG_INFO

Attribute	Description
TRANSACTION_LAG	**Deprecated**. To implement flow-of-control notifications, use the NTFN_GROUPING_* attributes. Applies only to OCN (described in "Object Change Notification (OCN)" on page 11-2). Has no effect if you specify the QOS_FLAGS attribute with its QOS_QUERY flag. Specifies the number of transactions or database changes by which the client can lag behind the database. If 0, then the client receives an invalidation message as soon as it is generated. If 5, then every fifth transaction that changes a registered object results in a notification. The database tracks intervening changes at an object granularity and bundles the changes along with the notification. Thus, the client does not lose intervening changes. Most applications that must be notified of changes to an object on transaction commit without further deferral are expected to chose 0 transaction lag. A nonzero transaction lag is useful only if an application implements flow control on notifications. When using nonzero transaction lag, it is recommended that the application workload has the property that notifications are generated at a reasonable frequency. Otherwise, notifications might be deferred indefinitely till the lag is satisfied. If you specify TRANSACTION_LAG, then the ROWID level granularity is not available in the notification messages even if you specified QOS_ROWIDS during registration.
NTFN_GROUPING_CLASS	Specifies the class by which to group notifications. The only allowed value is DBMS_CQ_NOTIFICATION.NTFN_GROUPING_CLASS_TIME, which groups notifications by time.
NTFN_GROUPING_VALUE	Specifies the time interval that defines the group, in seconds. For example, if this value is 900, notifications generated in the same 15-minute interval are grouped.
NTFN_GROUPING_TYPE	Specifies either of these types of grouping: ■ DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_SUMMARY: All notifications in the group are summarized into a single notification. ■ DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_LAST: Only the last notification in the group is published and the earlier ones discarded.
NTFN_GROUPING_START_TIME	Specifies when to start generating notifications. If specified as NULL, it defaults to the current system-generated time.
NTFN_GROUPING_REPEAT_COUNT	Specifies how many times to repeat the notification. Set to DBMS_CQ_NOTIFICATION.NTFN_GROUPING_FOREVER to receive notifications for the life of the registration. To receive at most *n* notifications during the life of the registration, set to *n*.

The quality-of-service flags in Table 11–3 are constants in the DBMS_CQ_NOTIFICATION package. You can specify them with the QOS_FLAGS attribute of CQ_NOTIFICATION$_REG_INFO (see Table 11–2).

Table 11–3 Quality-of-Service Flags

Flag	Description
QOS_DEREG_NFY	Purges the registration after the first notification.
QOS_RELIABLE	Stores notifications in a persistent database queue. In an Oracle RAC environment, if a database instance fails, surviving database instances can deliver any queued notification messages. **Default:** Notifications are stored in shared memory, which performs more efficiently.
QOS_ROWIDS	Includes the ROWID of each changed row in the notification.
QOS_QUERY	Registers queries for QRCN, described in Query Result Change Notification (QRCN) on page 11-2. If a query cannot be registered for QRCN, an error is generated at registration time, unless you also specify QOS_BEST_EFFORT. **Default:** Queries are registered for OCN, described in "Object Change Notification (OCN)" on page 11-2
QOS_BEST_EFFORT	Used with QOS_QUERY. Registers simplified versions of queries that are too complex for query result change evaluation; in other words, registers queries for QRCN in best-effort mode, described in "Best-Effort Mode" on page 11-3. To see which queries were simplified, query the static data dictionary view DBA_CQ_NOTIFICATION_QUERIES or USER_CQ_NOTIFICATION_QUERIES. These views give the QUERYID and the text of each registered query. **Default:** Queries are registered for QRCN in guaranteed mode, described in "Guaranteed Mode" on page 11-3

Suppose that you must invoke the procedure HR.dcn_callback whenever a registered object changes. In Example 11–4, you create a CQ_NOTIFICATION$_REG_INFO object that specifies that HR.dcn_callback receives notifications. To create the object you must have EXECUTE privileges on the DBMS_CQ_NOTIFICATION package.

Example 11–4 Creating a CQ_NOTIFICATION$_REG_INFO Object

```
DECLARE
  v_cn_addr CQ_NOTIFICATION$_REG_INFO;

BEGIN
  -- Create object:

  v_cn_addr := CQ_NOTIFICATION$_REG_INFO (
    'HR.dcn_callback',                 -- PL/SQL notification handler
    DBMS_CQ_NOTIFICATION.QOS_QUERY     -- notification type QRCN
    + DBMS_CQ_NOTIFICATION.QOS_ROWIDS, -- include rowids of changed objects
    0,                           -- registration persists until unregistered
    0,                           -- notify on all operations
    0                            -- notify immediately
    );

  -- Register queries: ...
END;
/
```

Identifying Individual Queries in a Notification

Any query in a registered list of queries can cause a continuous query notification. To know when a certain query causes a notification, use the DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID function in the SELECT list of that query. For example:

```
SELECT EMPLOYEE_ID, SALARY, DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID
FROM EMPLOYEES
WHERE DEPARTMENT_ID = 10;
```

Result:

```
EMPLOYEE_ID     SALARY CQ_NOTIFICATION_QUERYID
----------- ---------- -----------------------
        200       2800                       0

1 row selected.
```

When that query causes a notification, the notification includes the query ID.

Adding Queries to an Existing Registration

To add queries to an existing registration, follow these steps:

1. Retrieve the registration ID of the existing registration.

   You can retrieve it from either saved output or a query of *_CHANGE_NOTIFICATION_REGS.

2. Open the existing registration by calling the procedure DBMS_CQ_NOTIFICATION.ENABLE_REG with the registration ID as the parameter.
3. Run the queries to register. (Do not run DML or DDL operations.)
4. Close the registration, using the DBMS_CQ_NOTIFICATION.REG_END function.

Example 11–5 adds a query to an existing registration whose registration ID is 21.

Example 11–5 Adding a Query to an Existing Registration

```
DECLARE
  v_cursor SYS_REFCURSOR;

BEGIN
  -- Open existing registration
  DBMS_CQ_NOTIFICATION.ENABLE_REG(21);
  OPEN v_cursor FOR
    -- Run query to be registered
    SELECT DEPARTMENT_ID
      FROM HR.DEPARTMENTS;  -- register this query
  CLOSE v_cursor;
  -- Close registration
  DBMS_CQ_NOTIFICATION.REG_END;
END;
/
```

Best Practices for CQN Registrations

For best CQN performance, follow these registration guidelines:

- Register few queries—preferably those that reference objects that rarely change.

Extremely volatile registered objects cause numerous notifications, whose overhead slows OLTP throughput.

- Minimize the number of duplicate registrations of any given object, to avoid replicating a notification message for multiple recipients.

Troubleshooting CQN Registrations

If you are unable to create a registration, or if you have created a registration but are not receiving the notifications that you expected, the problem might be one of these:

- The `JOB_QUEUE_PROCESSES` parameter is not set to a nonzero value.

  This prevents you from receiving PL/SQL notifications through the notification handler.

- You were connected as a SYS user when you created the registrations.

  You must be connected as a non-SYS user to create CQN registrations.

- You changed a registered object, but did not commit the transaction.

  Notifications are generated only when the transaction commits.

- The registrations were not successfully created in the database.

  To check, query the static data dictionary view `*_CHANGE_NOTIFICATION_REGS`. For example, this statement displays all registrations and registered objects for the current user:

  ```
  SELECT REGID, TABLE_NAME FROM USER_CHANGE_NOTIFICATION_REGS;
  ```

- Run-time errors occurred during the execution of the notification handler.

  If so, they were logged to the trace file of the `JOBQ` process that tried to run the procedure. The name of the trace file usually has this form:

  ```
  ORACLE_SID_jnumber_PID.trc
  ```

  For example, if the ORACLE_SID is `dbs1` and the process ID (PID) of the `JOBQ` process is 12483, the name of the trace file is usually `dbs1_j000_12483.trc`.

  Suppose that a registration is created with `'chnf_callback'` as the notification handler and registration ID 100. Suppose that `'chnf_callback'` was not defined in the database. Then the `JOBQ` trace file might contain a message of the form:

  ```
  ****************************************************************************
     Run-time error during execution of PL/SQL cbk chnf_callback for reg CHNF100.
     Error in PLSQL notification of msgid:
     Queue :
     Consumer Name :
     PLSQL function :chnf_callback
     Exception Occured, Error msg:
     ORA-00604: error occurred at recursive SQL level 2
     ORA-06550: line 1, column 7:
     PLS-00201: identifier 'CHNF_CALLBACK' must be declared
     ORA-06550: line 1, column 7:
     PL/SQL: Statement ignored
  ****************************************************************************
  ```

  If run-time errors occurred during the execution of the notification handler, create a very simple version of the notification handler to verify that you are actually receiving notifications, and then gradually add application logic.

An example of a very simple notification handler is:

```
REM Create table in HR schema to hold count of notifications received.
CREATE TABLE nfcount(cnt NUMBER);
INSERT INTO nfcount (cnt) VALUES(0);
COMMIT;
CREATE OR REPLACE PROCEDURE chnf_callback
  (ntfnds IN CQ_NOTIFICATION$_DESCRIPTOR)
IS
BEGIN
  UPDATE nfcount SET cnt = cnt+1;
  COMMIT;
END;
/
```

- There is a time lag between the commit of a transaction and the notification received by the end user.

Querying CQN Registrations

To see top-level information about all registrations, including their QOS options, query the static data dictionary view `*_CHANGE_NOTIFICATION_REGS`.

For example, you can obtain the registration ID for a client and the list of objects for which it receives notifications. To view registration IDs and table names for `HR`, use this query:

```
SELECT regid, table_name FROM USER_CHANGE_NOTIFICATION_REGS;
```

To see which queries are registered for QRCN, query the static data dictionary view `USER_CQ_NOTIFICATION_QUERIES` or `DBA_CQ_NOTIFICATION_QUERIES`. These views include information about any bind values that the queries use. In these views, bind values in the original query are included in the query text as constants. The query text is equivalent, but maybe not identical, to the original query that was registered.

> **See Also:** *Oracle Database Reference* for more information about the static data dictionary views `USER_CHANGE_NOTIFICATION_REGS` and `DBA_CQ_NOTIFICATION_QUERIES`

Interpreting Notifications

When a transaction commits, the database determines whether registered objects were modified in the transaction. If so, it runs the notification handler specified in the registration.

Topics:

- Interpreting a CQ_NOTIFICATION$_DESCRIPTOR Object
- Interpreting a CQ_NOTIFICATION$_TABLE Object
- Interpreting a CQ_NOTIFICATION$_QUERY Object
- Interpreting a CQ_NOTIFICATION$_ROW Object

Interpreting a CQ_NOTIFICATION$_DESCRIPTOR Object

When a CQN registration generates a notification, the database passes a `CQ_NOTIFICATION$_DESCRIPTOR` object to the notification handler. The notification

handler can find the details of the database change in the attributes of the CQ_NOTIFICATION$_DESCRIPTOR object.

In SQL*Plus, you can list these attributes by connecting as SYS and running this statement:

```
DESC CQ_NOTIFICATION$_DESCRIPTOR
```

Table 11–4 summarizes the attributes of CQ_NOTIFICATION$_DESCRIPTOR.

Table 11–4 Attributes of CQ_NOTIFICATION$_DESCRIPTOR

Attribute	Description
REGISTRATION_ID	The registration ID that was returned during registration.
TRANSACTION_ID	The ID for the transaction that made the change.
DBNAME	The name of the database in which the notification was generated.
EVENT_TYPE	The database event that triggers a notification. For example, the attribute can contain these constants, which correspond to different database events: ■ EVENT_NONE ■ EVENT_STARTUP (Instance startup) ■ EVENT_SHUTDOWN (Instance shutdown - last instance shutdown for Oracle RAC) ■ EVENT_SHUTDOWN_ANY (Any instance shutdown for Oracle RAC) ■ EVENT_DEREG (Registration was removed) ■ EVENT_OBJCHANGE (Change to a registered table) ■ EVENT_QUERYCHANGE (Change to a registered result set)
NUMTABLES	The number of tables that were modified.
TABLE_DESC_ARRAY	This field is present only for OCN registrations. For QRCN registrations, it is NULL. If EVENT_TYPE is EVENT_OBJCHANGE]: a VARRAY of table change descriptors of type CQ_NOTIFICATION$_TABLE, each of which corresponds to a changed table. For attributes of CQ_NOTIFICATION$_TABLE, see Table 11–5. Otherwise: NULL.
QUERY_DESC_ARRAY	This field is present only for QRCN registrations. For OCN registrations, it is NULL. If EVENT_TYPE is EVENT_QUERYCHANGE]: a VARRAY of result set change descriptors of type CQ_NOTIFICATION$_QUERY, each of which corresponds to a changed result set. For attributes of CQ_NOTIFICATION$_QUERY, see Table 11–6. Otherwise: NULL.

Interpreting a CQ_NOTIFICATION$_TABLE Object

The CQ_NOTIFICATION$_DESCRIPTOR type contains an attribute called TABLE_DESC_ARRAY, which holds a VARRAY of table descriptors of type CQ_NOTIFICATION$_TABLE.

In SQL*Plus, you can list these attributes by connecting as SYS and running this statement:

```
DESC CQ_NOTIFICATION$_TABLE
```

Table 11–5 summarizes the attributes of CQ_NOTIFICATION$_TABLE.

Table 11–5 Attributes of CQ_NOTIFICATION$_TABLE

Attribute	Specifies . . .
OPFLAGS	The type of operation performed on the modified table. For example, the attribute can contain these constants, which correspond to different database operations: ■ ALL_ROWS signifies that either the entire table is modified, as in a DELETE *, or row-level granularity of information is not requested or not available in the notification, and the recipient must assume that the entire table has changed ■ UPDATEOP signifies an update ■ DELETEOP signifies a deletion ■ ALTEROP signifies an ALTER TABLE ■ DROPOP signifies a DROP TABLE ■ UNKNOWNOP signifies an unknown operation
TABLE_NAME	The name of the modified table.
NUMROWS	The number of modified rows.
ROW_DESC_ARRAY	A VARRAY of row descriptors of type CQ_NOTIFICATION$_ROW, which Table 11–7 describes. If ALL_ROWS was set in the opflags, then the ROW_DESC_ARRAY member is NULL.

Interpreting a CQ_NOTIFICATION$_QUERY Object

The CQ_NOTIFICATION$_DESCRIPTOR type contains an attribute called QUERY_DESC_ARRAY, which holds a VARRAY of result set change descriptors of type CQ_NOTIFICATION$_QUERY.

In SQL*Plus, you can list these attributes by connecting as SYS and running this statement:

```
DESC CQ_NOTIFICATION$_QUERY
```

Table 11–6 summarizes the attributes of CQ_NOTIFICATION$_QUERY.

Table 11–6 Attributes of CQ_NOTIFICATION$_QUERY

Attribute	Specifies . . .
QUERYID	Query ID of the changed query.
QUERYOP	Operation that changed the query (either EVENT_QUERYCHANGE or EVENT_DEREG).
TABLE_DESC_ARRAY	A VARRAY of table change descriptors of type CQ_NOTIFICATION$_TABLE, each of which corresponds to a changed table that caused a change in the result set. For attributes of CQ_NOTIFICATION$_TABLE, see Table 11–5.

Interpreting a CQ_NOTIFICATION$_ROW Object

If the ROWID option was specified during registration, the CQ_NOTIFICATION$_TABLE type has a ROW_DESC_ARRAY attribute, a VARRAY of type CQ_NOTIFICATION$_ROW that contains the ROWIDs for the changed rows. If ALL_ROWS was set in the OPFLAGS field of the CQ_NOTIFICATION$_TABLE object, then ROWID information is not available.

Table 11–7 summarizes the attributes of CQ_NOTIFICATION$_ROW.

Table 11–7 Attributes of CQ_NOTIFICATION$_ROW

Attribute	Specifies . . .
OPFLAGS	The type of operation performed on the modified table. See the description of OPFLAGS in Table 11–5.
ROW_ID	The ROWID of the changed row.

Deleting Registrations

To delete a registration, call the procedure DBMS_CQ_NOTIFICATION.DEREGISTER with the registration ID as the parameter. For example, this statement deregisters the registration whose registration ID is 21:

```
DBMS_CQ_NOTIFICATION.DEREGISTER(21);
```

Only the user who created the registration or the SYS user can deregister it.

Configuring CQN: Scenario

In this scenario, you are a developer who manages a Web application that provides employee data: name, location, phone number, and so on. The application, which runs on Oracle Application Server, is heavily used and processes frequent queries of the HR.EMPLOYEES and HR.DEPARTMENTS tables in the back-end database. Because these tables change relatively infrequently, the application can improve performance by caching the query results. Caching avoids a round trip to the back-end database and server-side execution latency.

You can use the DBMS_CQ_NOTIFICATION package to register queries based on HR.EMPLOYEES and HR.DEPARTMENTS tables. To configure CQN, you follow these steps:

1. Create a server-side PL/SQL stored procedure to process the notifications, as instructed in "Creating a PL/SQL Notification Handler" on page 11-27.
2. Register the queries on the HR.EMPLOYEES and HR.DEPARTMENTS tables for QRCN, as instructed in "Registering the Queries" on page 11-29.

After you complete these steps, any committed change to the result of a query registered in step 2 causes the notification handler created in step 1 to notify the Web application of the change, whereupon the Web application refreshes the cache by querying the back-end database.

Topics:

- Creating a PL/SQL Notification Handler
- Registering the Queries

Creating a PL/SQL Notification Handler

Create a a server-side stored PL/SQL procedure to process notifications as follows:

1. Connect to the database AS SYSDBA.
2. Grant the required privileges to HR:

```
GRANT EXECUTE ON DBMS_CQ_NOTIFICATION TO HR;
GRANT CHANGE NOTIFICATION TO HR;
```

3. Enable the JOB_QUEUE_PROCESSES parameter to receive notifications:

```
ALTER SYSTEM SET "JOB_QUEUE_PROCESSES"=4;
```

4. Connect to the database as a non-SYS user (such as HR).
5. Create database tables to hold records of notification events received:

```
-- Create table to record notification events.
DROP TABLE nfevents;
CREATE TABLE nfevents (
  regid      NUMBER,
  event_type NUMBER
);

-- Create table to record notification queries:
DROP TABLE nfqueries;
CREATE TABLE nfqueries (
  qid NUMBER,
  qop NUMBER
);

-- Create table to record changes to registered tables:
DROP TABLE nftablechanges;
CREATE TABLE nftablechanges (
  qid             NUMBER,
  table_name      VARCHAR2(100),
  table_operation NUMBER
);

-- Create table to record ROWIDs of changed rows:
DROP TABLE nfrowchanges;
CREATE TABLE nfrowchanges (
  qid        NUMBER,
  table_name VARCHAR2(100),
  row_id     VARCHAR2(2000)
);
```

6. Create the procedure HR.chnf_callback, as shown in Example 11–6.

Example 11–6 Creating Server-Side PL/SQL Notification Handler

```
CREATE OR REPLACE PROCEDURE chnf_callback (
  ntfnds IN CQ_NOTIFICATION$_DESCRIPTOR
)
IS
  regid           NUMBER;
  tbname          VARCHAR2(60);
  event_type      NUMBER;
  numtables       NUMBER;
  operation_type  NUMBER;
  numrows         NUMBER;
  row_id          VARCHAR2(2000);
  numqueries      NUMBER;
  qid             NUMBER;
  qop             NUMBER;

BEGIN
  regid := ntfnds.registration_id;
  event_type := ntfnds.event_type;

  INSERT INTO nfevents (regid, event_type)
  VALUES (chnf_callback.regid, chnf_callback.event_type);
```

```
  numqueries :=0;

  IF (event_type = DBMS_CQ_NOTIFICATION.EVENT_QUERYCHANGE) THEN
    numqueries := ntfnds.query_desc_array.count;

    FOR i IN 1..numqueries LOOP  -- loop over queries
      qid := ntfnds.query_desc_array(i).queryid;
      qop := ntfnds.query_desc_array(i).queryop;

      INSERT INTO nfqueries (qid, qop)
      VALUES(chnf_callback.qid, chnf_callback.qop);

      numtables := 0;
      numtables := ntfnds.query_desc_array(i).table_desc_array.count;

      FOR j IN 1..numtables LOOP  -- loop over tables
        tbname :=
          ntfnds.query_desc_array(i).table_desc_array(j).table_name;
        operation_type :=
          ntfnds.query_desc_array(i).table_desc_array(j).Opflags;

        INSERT INTO nftablechanges (qid, table_name, table_operation)
        VALUES (
          chnf_callback.qid,
          tbname,
          operation_type
        );

        IF (bitand(operation_type, DBMS_CQ_NOTIFICATION.ALL_ROWS) = 0) THEN
          numrows := ntfnds.query_desc_array(i).table_desc_array(j).numrows;
        ELSE
          numrows :=0;  -- ROWID info not available
        END IF;

        -- Body of loop does not run when numrows is zero.
        FOR k IN 1..numrows LOOP  -- loop over rows
          Row_id :=
 ntfnds.query_desc_array(i).table_desc_array(j).row_desc_array(k).row_id;

          INSERT INTO nfrowchanges (qid, table_name, row_id)
          VALUES (chnf_callback.qid, tbname, chnf_callback.Row_id);

        END LOOP;  -- loop over rows
      END LOOP;  -- loop over tables
    END LOOP;  -- loop over queries
  END IF;
  COMMIT;
END;
/
```

Registering the Queries

After creating the notification handler, you register the queries for which you want to receive notifications, specifying `HR.chnf_callback` as the notification handler, as in Example 11–7.

Example 11–7 Registering a Query

```
DECLARE
```

```
  reginfo   CQ_NOTIFICATION$_REG_INFO;
  mgr_id    NUMBER;
  dept_id   NUMBER;
  v_cursor  SYS_REFCURSOR;
  regid     NUMBER;

BEGIN
  /* Register two queries for QRNC: */
  /* 1. Construct registration information.
        chnf_callback is name of notification handler.
        QOS_QUERY specifies result-set-change notifications. */

  reginfo := cq_notification$_reg_info (
    'chnf_callback',
    DBMS_CQ_NOTIFICATION.QOS_QUERY,
    0, 0, 0
  );

  /* 2. Create registration. */

  regid := DBMS_CQ_NOTIFICATION.new_reg_start(reginfo);

  OPEN v_cursor FOR
    SELECT dbms_cq_notification.CQ_NOTIFICATION_QUERYID, manager_id
    FROM HR.EMPLOYEES
    WHERE employee_id = 7902;
  CLOSE v_cursor;

  OPEN v_cursor FOR
    SELECT dbms_cq_notification.CQ_NOTIFICATION_QUERYID, department_id
    FROM HR.departments
    WHERE department_name = 'IT';
  CLOSE v_cursor;

  DBMS_CQ_NOTIFICATION.reg_end;
END;
/
```

View the newly created registration:

```
SELECT queryid, regid, TO_CHAR(querytext)
FROM user_cq_notification_queries;
```

Result is similar to:

```
QUERYID REGID                              TO_CHAR(QUERYTEXT)
------- ----- ------------------------------------------------
     22    41 SELECT HR.DEPARTMENTS.DEPARTMENT_ID
                FROM HR.DEPARTMENTS
                  WHERE HR.DEPARTMENTS.DEPARTMENT_NAME  = 'IT'

     21    41 SELECT HR.EMPLOYEES.MANAGER_ID
                FROM HR.EMPLOYEES
                  WHERE HR.EMPLOYEES.EMPLOYEE_ID  = 7902
```

Run this transaction, which changes the result of the query with `QUERYID` 22:

```
UPDATE DEPARTMENTS
SET DEPARTMENT_NAME = 'FINANCE'
WHERE department_name = 'IT';
```

The notification procedure `chnf_callback` (which you created in Example 11–6) runs.

Query the table in which notification events are recorded:

```
SELECT * FROM nfevents;
```

Result is similar to:

```
REGID EVENT_TYPE
----- ----------
   61          7
```

`EVENT_TYPE` 7 corresponds to `EVENT_QUERYCHANGE` (query result change).

Query the table in which changes to registered tables are recorded:

```
SELECT * FROM nftablechanges;
```

Result is similar to:

```
REGID     TABLE_NAME TABLE_OPERATION
----- -------------- ---------------
   42 HR.DEPARTMENTS               4
```

`TABLE_OPERATION` 4 corresponds to `UPDATEOP` (update operation).

Query the table in which `ROWID`s of changed rows are recorded:

```
SELECT * FROM nfrowchanges;
```

Result is similar to:

```
REGID     TABLE_NAME              ROWID
----- -------------- ------------------
   61 HR.DEPARTMENTS AAANkdAABAAALinAAF
```

Part III

Advanced Topics for Application Developers

This part presents application development information that either involves sophisticated technology or is used by a small minority of developers.

Chapters:

- Chapter 12, "Using Oracle Flashback Technology"
- Chapter 13, "Choosing a Programming Environment"
- Chapter 14, "Developing Applications with Multiple Programming Languages"
- Chapter 15, "Developing Applications with Oracle XA"
- Chapter 16, "Developing Applications with the Publish-Subscribe Model"
- Chapter 17, "Using the Identity Code Package"
- Chapter 18, "Schema Object Dependency"
- Chapter 19, "Edition-Based Redefinition"

> **See Also:** *Oracle Database Performance Tuning Guide* for performance issues to consider when developing applications

12

Using Oracle Flashback Technology

This chapter explains how to use Oracle Flashback Technology in database applications.

Topics:

- Overview of Oracle Flashback Technology
- Configuring Your Database for Oracle Flashback Technology
- Using Oracle Flashback Query (SELECT AS OF)
- Using Oracle Flashback Version Query
- Using Oracle Flashback Transaction Query
- Using Oracle Flashback Transaction Query with Oracle Flashback Version Query
- Using ORA_ROWSCN
- Using DBMS_FLASHBACK Package
- Using Flashback Transaction
- Using Flashback Data Archive (Oracle Total Recall)
- General Guidelines for Oracle Flashback Technology
- Performance Guidelines for Oracle Flashback Technology

Overview of Oracle Flashback Technology

Oracle Flashback Technology is a group of Oracle Database features that let you view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.

With flashback features, you can:

- Perform queries that return past data
- Perform queries that return metadata that shows a detailed history of changes to the database
- Recover tables or rows to a previous point in time
- Automatically track and archive transactional data changes
- Roll back a transaction and its dependent transactions while the database remains online

Oracle Flashback features use the Automatic Undo Management (AUM) system to obtain metadata and historical data for transactions. They rely on **undo data**, which

are records of the effects of individual transactions. For example, if a user runs an `UPDATE` statement to change a salary from 1000 to 1100, then Oracle Database stores the value 1000 in the undo data.

Undo data is persistent and survives a database shutdown. By using flashback features, you can use undo data to query past data or recover from logical damage. Besides using it in flashback features, Oracle Database uses undo data to perform these actions:

- Roll back active transactions
- Recover terminated transactions by using database or process recovery
- Provide read consistency for SQL queries

Topics:

- Application Development Features
- Database Administration Features

For additional general information about flashback features, see *Oracle Database Concepts*

Application Development Features

In application development, you can use these flashback features to report historical data or undo erroneous changes. (You can also use these features interactively as a database user or administrator.)

Oracle Flashback Query

Use this feature to retrieve data for an earlier time that you specify with the `AS OF` clause of the `SELECT` statement. For more information, see "Using Oracle Flashback Query (SELECT AS OF)" on page 12-6.

Oracle Flashback Version Query

Use this feature to retrieve metadata and historical data for a specific time interval (for example, to view all the rows of a table that ever existed during a given time interval). Metadata for each row version includes start and end time, type of change operation, and identity of the transaction that created the row version. To create an Oracle Flashback Version Query, use the `VERSIONS BETWEEN` clause of the `SELECT` statement. For more information, see "Using Oracle Flashback Version Query" on page 12-8.

Oracle Flashback Transaction Query

Use this feature to retrieve metadata and historical data for a given transaction or for all transactions in a given time interval. To perform an Oracle Flashback Transaction Query, select from the static data dictionary view `FLASHBACK_TRANSACTION_QUERY`. For more information, see "Using Oracle Flashback Transaction Query" on page 12-9.

Typically, you use Oracle Flashback Transaction Query with an Oracle Flashback Version Query that provides the transaction IDs for the rows of interest (see "Using Oracle Flashback Transaction Query with Oracle Flashback Version Query" on page 12-10).

DBMS_FLASHBACK Package

Use this feature to set the internal Oracle Database clock to an earlier time so that you can examine data that was current at that time, or to roll back a transaction and its

dependent transactions while the database remains online (see Flashback Transaction). For more information, see "Using DBMS_FLASHBACK Package" on page 12-14.

Flashback Transaction

Use Flashback Transaction to roll back a transaction and its dependent transactions while the database remains online. This recovery operation uses undo data to create and run the corresponding compensating transactions that return the affected data to its original state. (Flashback Transaction is part of `DBMS_FLASHBACK` package.) For more information, see "Using DBMS_FLASHBACK Package" on page 12-14.

Flashback Data Archive (Oracle Total Recall)

Use Flashback Data Archive to automatically track and archive both regular queries and Oracle Flashback Query, ensuring SQL-level access to the versions of database objects without getting a snapshot-too-old error. For more information, see "Using Flashback Data Archive (Oracle Total Recall)" on page 12-18.

Database Administration Features

These flashback features are primarily for data recovery. Typically, you use these features only as a database administrator.

This chapter focuses on the "Application Development Features" on page 12-2. For more information about the database administration features, see *Oracle Database Administrator's Guide* and the *Oracle Database Backup and Recovery User's Guide*.

Oracle Flashback Table

Use this feature to restore a table to its state at a previous point in time. You can restore a table while the database is on line, undoing changes to only the specified table.

Oracle Flashback Drop

Use this feature to recover a dropped table. This feature reverses the effects of a `DROP TABLE` statement.

Oracle Flashback Database

Use this feature to quickly return the database to an earlier point in time, by undoing all of the changes that have taken place since then. This is fast, because you do not have to restore database backups.

Configuring Your Database for Oracle Flashback Technology

Before you can use flashback features in your application, you or your database administrator must perform the configuration tasks described in these topics:

- Configuring Your Database for Automatic Undo Management
- Configuring Your Database for Oracle Flashback Transaction Query
- Configuring Your Database for Flashback Transaction
- Enabling Oracle Flashback Operations on Specific LOB Columns
- Granting Necessary Privileges

Configuring Your Database for Automatic Undo Management

To configure your database for Automatic Undo Management (AUM), you or your database administrator must:

- Create an undo tablespace with enough space to keep the required data for flashback operations.

  The more often users update the data, the more space is required. The database administrator usually calculates the space requirement.

- Enable AUM, as explained in *Oracle Database Administrator's Guide*. Set these database initialization parameters:

  – `UNDO_MANAGEMENT`

  – `UNDO_TABLESPACE`

  – `UNDO_RETENTION`

  For a fixed-size undo tablespace, Oracle Database automatically tunes the system to give the undo tablespace the best possible undo retention.

  For an automatically extensible undo tablespace, Oracle Database retains undo data longer than the longest query duration and the low threshold of undo retention specified by the `UNDO_RETENTION` parameter.

  > **Note:** You can query `V$UNDOSTAT.TUNED_UNDORETENTION` to determine the amount of time for which undo is retained for the current undo tablespace. For more information about `V$UNDOSTAT`, see *Oracle Database Reference*.

  Setting `UNDO_RETENTION` does not guarantee that unexpired undo data is not discarded. If the system needs more space, Oracle Database can overwrite unexpired undo with more recently generated undo data.

- Specify the `RETENTION GUARANTEE` clause for the undo tablespace to ensure that unexpired undo data is not discarded.

  > **See Also:** *Oracle Database Administrator's Guide* for more information about creating an undo tablespace and enabling AUM

Configuring Your Database for Oracle Flashback Transaction Query

To configure your database for the Oracle Flashback Transaction Query feature, you or your database administrator must:

- Ensure that Oracle Database is running with version 10.0 compatibility.
- Enable supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Configuring Your Database for Flashback Transaction

To configure your database for the Flashback Transaction feature, you or your database administrator must:

- With the database mounted but not open, enable `ARCHIVELOG`:

```
ALTER DATABASE ARCHIVELOG;
```

- Open at least one archive log:

```
ALTER SYSTEM ARCHIVE LOG CURRENT;
```

- If not done, enable minimal and primary key supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (PRIMARY KEY) COLUMNS;
```

- If you want to track foreign key dependencies, enable foreign key supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (FOREIGN KEY) COLUMNS;
```

> **Note:** If you have very many foreign key constraints, enabling foreign key supplemental logging might not be worth the performance penalty.

Enabling Oracle Flashback Operations on Specific LOB Columns

To enable flashback operations on specific LOB columns of a table, use the `ALTER TABLE` statement with the `RETENTION` option.

Because undo data for LOB columns can be voluminous, you must define which LOB columns to use with flashback operations.

> **See Also:** *Oracle Database SecureFiles and Large Objects Developer's Guide* to learn about LOB storage and the `RETENTION` parameter

Granting Necessary Privileges

You or your database administrator must grant privileges to users, roles, or applications that must use these flashback features. For information about the `GRANT` statement, see *Oracle Database SQL Language Reference*.

For Oracle Flashback Query and Oracle Flashback Version Query

To allow access to specific objects during queries, grant `FLASHBACK` and `SELECT` privileges on those objects.

To allow queries on all tables, grant the `FLASHBACK ANY TABLE` privilege.

For Oracle Flashback Transaction Query

Grant the `SELECT ANY TRANSACTION` privilege.

To allow execution of undo SQL code retrieved by an Oracle Flashback Transaction Query, grant `SELECT`, `UPDATE`, `DELETE`, and `INSERT` privileges for specific tables.

For DBMS_FLASHBACK Package

To allow access to the features in the `DBMS_FLASHBACK` package, grant the `EXECUTE` privilege on `DBMS_FLASHBACK`.

For Flashback Data Archive (Oracle Total Recall)

To allow a specific user to enable Flashback Data Archive on tables, using a specific Flashback Data Archive, grant the `FLASHBACK ARCHIVE` object privilege on that Flashback Data Archive to that user. To grant the `FLASHBACK ARCHIVE` object privilege, you must either be logged on as `SYSDBA` or have `FLASHBACK ARCHIVE ADMINISTER` system privilege.

To allow execution of these statements, grant the `FLASHBACK ARCHIVE ADMINISTER` system privilege:

- `CREATE FLASHBACK ARCHIVE`
- `ALTER FLASHBACK ARCHIVE`
- `DROP FLASHBACK ARCHIVE`

To grant the `FLASHBACK ARCHIVE ADMINISTER` system privilege, you must be logged on as `SYSDBA`.

To create a default Flashback Data Archive, using either the `CREATE FLASHBACK ARCHIVE` or `ALTER FLASHBACK ARCHIVE` statement, you must be logged on as `SYSDBA`.

To disable Flashback Data Archive for a table that has been enabled for Flashback Data Archive, you must either be logged on as `SYSDBA` or have the `FLASHBACK ARCHIVE ADMINISTER` system privilege.

Using Oracle Flashback Query (SELECT AS OF)

To use Oracle Flashback Query, use a `SELECT` statement with an `AS OF` clause. Oracle Flashback Query retrieves data as it existed at an earlier time. The query explicitly references a past time through a time stamp or System Change Number (SCN). It returns committed data that was current at that point in time.

Uses of Oracle Flashback Query include:

- Recovering lost data or undoing incorrect, committed changes.

  For example, if you mistakenly delete or update rows, and then commit them, you can immediately undo the mistake.
- Comparing current data with the corresponding data at an earlier time.

  For example, you can run a daily report that shows the change in data from yesterday. You can compare individual rows of table data or find intersections or unions of sets of rows.
- Checking the state of transactional data at a particular time.

  For example, you can verify the account balance of a certain day.
- Simplifying application design by removing the need to store some kinds of temporal data.

  Oracle Flashback Query lets you retrieve past data directly from the database.
- Applying packaged applications, such as report generation tools, to past data.
- Providing self-service error correction for an application, thereby enabling users to undo and correct their errors.

Topics:

- Example of Examining and Restoring Past Data
- Guidelines for Oracle Flashback Query

For more information about the `SELECT AS OF` statement, see *Oracle Database SQL Language Reference*.

Example of Examining and Restoring Past Data

Suppose that you discover at 12:30 PM that the row for employee Chung was deleted from the `employees` table, and you know that at 9:30AM the data for Chung was correctly stored in the database. You can use Oracle Flashback Query to examine the contents of the table at 9:30 AM to find out what data was lost. If appropriate, you can restore the lost data.

Example 12–1 retrieves the state of the record for `Chung` at 9:30AM, April 4, 2004:

Example 12–1 Retrieving a Lost Row with Oracle Flashback Query

```
SELECT * FROM employees
AS OF TIMESTAMP
TO_TIMESTAMP('2004-04-04 09:30:00', 'YYYY-MM-DD HH:MI:SS')
WHERE last_name = 'Chung';
```

Example 12–2 restores Chung's information to the `employees` table:

Example 12–2 Restoring a Lost Row After Oracle Flashback Query

```
INSERT INTO employees (
  SELECT * FROM employees
  AS OF TIMESTAMP
  TO_TIMESTAMP('2004-04-04 09:30:00', 'YYYY-MM-DD HH:MI:SS')
  WHERE last_name = 'Chung'
);
```

Guidelines for Oracle Flashback Query

- You can specify or omit the `AS OF` clause for each table and specify different times for different tables.

  Note: If a table is a Flashback Data Archive and you specify a time for it that is earlier than its creation time, the query returns zero rows for that table, rather than causing an error. (For information about Flashback Data Archives, see "Using Flashback Data Archive (Oracle Total Recall)" on page 12-18.)

- You can use the `AS OF` clause in queries to perform data definition language (DDL) operations (such as creating and truncating tables) or data manipulation language (DML) statements (such as `INSERT` and `DELETE`) in the same session as Oracle Flashback Query.
- To use the result of Oracle Flashback Query in a DDL or DML statement that affects the current state of the database, use an `AS OF` clause inside an `INSERT` or `CREATE TABLE AS SELECT` statement.
- If a possible 3-second error (maximum) is important to Oracle Flashback Query in your application, use an SCN instead of a time stamp. See "General Guidelines for Oracle Flashback Technology" on page 12-25.
- You can create a view that refers to past data by using the `AS OF` clause in the `SELECT` statement that defines the view.

  If you specify a relative time by subtracting from the current time on the database host, the past time is recalculated for each query. For example:

```
CREATE VIEW hour_ago AS
  SELECT * FROM employees
    AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '60' MINUTE);
```

SYSTIMESTAMP refers to the time zone of the database host environment.

- You can use the AS OF clause in self-joins, or in set operations such as INTERSECT and MINUS, to extract or compare data from two different times.

  You can store the results by preceding Oracle Flashback Query with a CREATE TABLE AS SELECT or INSERT INTO TABLE SELECT statement. For example, this query reinserts into table employees the rows that existed an hour ago:

```
INSERT INTO employees
    (SELECT * FROM employees
     AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '60' MINUTE)
    )
    MINUS SELECT * FROM employees);
```

  SYSTIMESTAMP refers to the time zone of the database host environment.

Using Oracle Flashback Version Query

Use Oracle Flashback Version Query to retrieve the different versions of specific rows that existed during a given time interval. A row version is created whenever a COMMIT statement is executed.

Specify Oracle Flashback Version Query using the VERSIONS BETWEEN clause of the SELECT statement. The syntax is:

```
VERSIONS {BETWEEN {SCN | TIMESTAMP} start AND end}
```

where *start* and *end* are expressions representing the start and end, respectively, of the time interval to be queried. The time interval includes (*start* and *end*).

Oracle Flashback Version Query returns a table with a row for each version of the row that existed at any time during the specified time interval. Each row in the table includes pseudocolumns of metadata about the row version, described in Table 12–1. This information can reveal when and how a particular change (perhaps erroneous) occurred to your database.

Table 12–1 Oracle Flashback Version Query Row Data Pseudocolumns

Pseudocolumn Name	Description
VERSIONS_STARTSCN VERSIONS_STARTTIME	Starting System Change Number (SCN) or TIMESTAMP when the row version was created. This pseudocolumn identifies the time when the data first had the values reflected in the row version. Use this pseudocolumn to identify the past target time for Oracle Flashback Table or Oracle Flashback Query. If this pseudocolumn is NULL, then the row version was created before *start*.
VERSIONS_ENDSCN VERSIONS_ENDTIME	SCN or TIMESTAMP when the row version expired. If this pseudocolumn is NULL, then either the row version was current at the time of the query or the row corresponds to a DELETE operation.
VERSIONS_XID	Identifier of the transaction that created the row version.

Table 12–1 (Cont.) Oracle Flashback Version Query Row Data Pseudocolumns

Pseudocolumn Name	Description
VERSIONS_OPERATION	Operation performed by the transaction: I for insertion, D for deletion, or U for update. The version is that of the row that was inserted, deleted, or updated; that is, the row after an INSERT operation, the row before a DELETE operation, or the row affected by an UPDATE operation. For user updates of an index key, Oracle Flashback Version Query might treat an UPDATE operation as two operations, DELETE plus INSERT, represented as two version rows with a D followed by an I VERSIONS_OPERATION.

A given row version is valid starting at its time VERSIONS_START* up to, but not including, its time VERSIONS_END*. That is, it is valid for any time *t* such that VERSIONS_START* <= *t* < VERSIONS_END*. For example, this output indicates that the salary was 10243 from September 9, 2002, included, to November 25, 2003, excluded.

```
VERSIONS_START_TIME     VERSIONS_END_TIME     SALARY
-------------------     -----------------     ------
09-SEP-2003             25-NOV-2003           10243
```

Here is a typical use of Oracle Flashback Version Query:

```
SELECT versions_startscn, versions_starttime,
       versions_endscn, versions_endtime,
       versions_xid, versions_operation,
       last_name, salary
  FROM employees
  VERSIONS BETWEEN TIMESTAMP
      TO_TIMESTAMP('2008-12-18 14:00:00', 'YYYY-MM-DD HH24:MI:SS')
  AND TO_TIMESTAMP('2008-12-18 17:00:00', 'YYYY-MM-DD HH24:MI:SS')
  WHERE first_name = 'John';
```

You can use VERSIONS_XID with Oracle Flashback Transaction Query to locate this transaction's metadata, including the SQL required to undo the row change and the user responsible for the change—see "Using Oracle Flashback Transaction Query" on page 12-9.

> **See Also:** *Oracle Database SQL Language Reference* for information about Oracle Flashback Version Query pseudocolumns and the syntax of the VERSIONS clause

Using Oracle Flashback Transaction Query

Use Oracle Flashback Transaction Query to retrieve metadata and historical data for a given transaction or for all transactions in a given time interval. Oracle Flashback Transaction Query queries the static data dictionary view FLASHBACK_TRANSACTION_QUERY, whose columns are described in *Oracle Database Reference*.

The column UNDO_SQL shows the SQL code that is the is the logical opposite of the DML operation performed by the transaction. You can usually use this code to reverse the logical steps taken during the transaction. However, there are cases where the SQL_UNDO code is not the exact opposite of the original transaction. For example, a SQL_UNDO INSERT operation might not insert a row back in a table at the same ROWID from which it was deleted.

This statement queries the FLASHBACK_TRANSACTION_QUERY view for transaction information, including the transaction ID, the operation, the operation start and end SCNs, the user responsible for the operation, and the SQL code that shows the logical opposite of the operation:

```
SELECT xid, operation, start_scn, commit_scn, logon_user, undo_sql
FROM flashback_transaction_query
WHERE xid = HEXTORAW('000200030000002D');
```

This statement uses Oracle Flashback Version Query as a subquery to associate each row version with the LOGON_USER responsible for the row data change:

```
SELECT xid, logon_user
FROM flashback_transaction_query
WHERE xid IN (
  SELECT versions_xid FROM employees VERSIONS BETWEEN TIMESTAMP
  TO_TIMESTAMP('2003-07-18 14:00:00', 'YYYY-MM-DD HH24:MI:SS') AND
  TO_TIMESTAMP('2003-07-18 17:00:00', 'YYYY-MM-DD HH24:MI:SS')
);
```

Note: If you query FLASHBACK_TRANSACTION_QUERY without specifying XID in the WHERE clause, the query scans many unrelated rows, degrading performance.

See Also:

- *Oracle Database Backup and Recovery User's Guide.* for information about how a database administrator can use Flashback Table to restore an entire table, rather than individual rows
- *Oracle Database Administrator's Guide* for information about how a database administrator can use Flashback Table to restore an entire table, rather than individual rows

Using Oracle Flashback Transaction Query with Oracle Flashback Version Query

In this example, a database administrator does this:

```
DROP TABLE emp;
CREATE TABLE emp (
  empno   NUMBER PRIMARY KEY,
  empname VARCHAR2(16),
  salary  NUMBER
);
INSERT INTO emp (empno, empname, salary) VALUES (111, 'Mike', 555);
COMMIT;

DROP TABLE dept;
CREATE TABLE dept (
  deptno   NUMBER,
  deptname VARCHAR2(32)
);
INSERT INTO dept (deptno, deptname) VALUES (10, 'Accounting');
COMMIT;
```

Now emp and dept have one row each. In terms of row versions, each table has one version of one row. Suppose that an erroneous transaction deletes empno 111 from table emp:

```
UPDATE emp SET salary = salary + 100 WHERE empno = 111;
INSERT INTO dept (deptno, deptname) VALUES (20, 'Finance');
DELETE FROM emp WHERE empno = 111;
COMMIT;
```

Next, a transaction reinserts empno 111 into the emp table with a new employee name:

```
INSERT INTO emp (empno, empname, salary) VALUES (111, 'Tom', 777);
UPDATE emp SET salary = salary + 100 WHERE empno = 111;
UPDATE emp SET salary = salary + 50 WHERE empno = 111;
COMMIT;
```

The database administrator detects the application error and must diagnose the problem. The database administrator issues this query to retrieve versions of the rows in the emp table that correspond to empno 111. The query uses Oracle Flashback Version Query pseudocolumns:

```
SELECT versions_xid XID, versions_startscn START_SCN,
  versions_endscn END_SCN, versions_operation OPERATION,
  empname, salary
FROM emp
VERSIONS BETWEEN SCN MINVALUE AND MAXVALUE
WHERE empno = 111;
```

Results are similar to:

```
XID               START_SCN    END_SCN O EMPNAME              SALARY
---------------- ---------- ---------- - ---------------- ----------
09001100B2200000   10093466            I Tom                     927
030002002B210000   10093459            D Mike                    555
0800120096200000   10093375   10093459 I Mike                    555

3 rows selected.
```

The results table rows are in descending chronological order. The third row corresponds to the version of the row in the table emp that was inserted in the table when the table was created. The second row corresponds to the row in emp that the erroneous transaction deleted. The first row corresponds to the version of the row in emp that was reinserted with a new employee name.

The database administrator identifies transaction 030002002B210000 as the erroneous transaction and uses Oracle Flashback Transaction Query to audit all changes made by this transaction:

```
SELECT  xid, start_scn, commit_scn, operation, logon_user, undo_sql
FROM flashback_transaction_query
WHERE xid = HEXTORAW('000200030000002D');
```

Results are similar to:

```
XID               START_SCN COMMIT_SCN OPERATION LOGON_USER
---------------- ---------- ---------- --------- ------------------------------
UNDO_SQL
--------------------------------------------------------------------------------

030002002B210000   10093452   10093459 DELETE    HR
insert into "HR"."EMP"("EMPNO","EMPNAME","SALARY") values ('111','Mike','655');
```

```
030002002B210000   10093452   10093459 INSERT    HR
delete from "HR"."DEPT" where ROWID = 'AAATjuAAEAAAAJrAAB';

030002002B210000   10093452   10093459 UPDATE    HR
update "HR"."EMP" set "SALARY" = '555' where ROWID = 'AAATjsAAEAAAAJ7AAA';

030002002B210000   10093452   10093459 BEGIN     HR


4 rows selected.
```

To make the result of the next query easier to read, the database administrator uses these SQL*Plus commands:

```
COLUMN operation FORMAT A9
COLUMN table_name FORMAT A10
COLUMN table_owner FORMAT A11
```

To see the details of the erroneous transaction and all subsequent transactions, the database administrator performs this query:

```
SELECT xid, start_scn, commit_scn, operation, table_name, table_owner
FROM flashback_transaction_query
WHERE table_owner = 'HR'
AND start_timestamp >=
  TO_TIMESTAMP ('2002-04-16 11:00:00','YYYY-MM-DD HH:MI:SS');
```

Results are similar to:

```
XID               START_SCN COMMIT_SCN OPERATION TABLE_NAME TABLE_OWNER
---------------- ---------- ---------- --------- ---------- -----------
02000E0074200000   10093435   10093446 INSERT    DEPT       HR
030002002B210000   10093452   10093459 DELETE    EMP        HR
030002002B210000   10093452   10093459 INSERT    DEPT       HR
030002002B210000   10093452   10093459 UPDATE    EMP        HR
0800120096200000   10093374   10093375 INSERT    EMP        HR
09001100B2200000   10093462   10093466 UPDATE    EMP        HR
09001100B2200000   10093462   10093466 UPDATE    EMP        HR
09001100B2200000   10093462   10093466 INSERT    EMP        HR

8 rows selected.
```

Note: Because the preceding query does not specify the `XID` in the `WHERE` clause, it scans many unrelated rows, degrading performance.

Using ORA_ROWSCN

`ORA_ROWSCN` is a pseudocolumn of any table that is not fixed or external. It represents the SCN of the most recent change to a given row in the current session; that is, the most recent `COMMIT` operation for the row in the current session. For example:

```
SELECT ora_rowscn, last_name, salary
FROM employees
WHERE employee_id = 200;
```

Result is similar to:

```
ORA_ROWSCN LAST_NAME                     SALARY
---------- ------------------------- ----------
```

```
884320 Whalen                          2800
```

The most recent COMMIT operation for the row in the current session took place at approximately SCN 9371092. To convert an SCN to the corresponding TIMESTAMP value, use the function SCN_TO_TIMESTAMP (documented in *Oracle Database SQL Language Reference*).

ORA_ROWSCN is a conservative upper bound of the latest commit time—the actual commit SCN can be somewhat earlier. ORA_ROWSCN is more precise (closer to the actual commit SCN) for a row-dependent table (created using CREATE TABLE with the ROWDEPENDENCIES clause). For more information about ORA_ROWSCN and ROWDEPENDENCIES, see *Oracle Database SQL Language Reference*.

Note: ORA_ROWSCN is not supported for Flashback Query. Instead, use the version query pseudocolumns, which are provided explicitly for Flashback Query. For information about these pseudocolumns, see *Oracle Database SQL Language Reference*.

Topics:

- Scenario: Package Subprogram Might Change Row
- ORA_ROWSCN and Tables with Virtual Private Database (VPD)

Scenario: Package Subprogram Might Change Row

Your application examines a row of data and records the corresponding ORA_ROWSCN as 202553. Then, your application invokes a package subprogram, whose implementation details you cannot see, which might or might not change the same row (and commit the change). Later, your application must update the row only if the package subprogram did not change it. Make the operation conditional—update the row only if ORA_ROWSCN is still 202553, as in this equivalent interactive statement:

```
UPDATE employees
SET salary = salary + 100
WHERE employee_id = 200
AND ora_rowscn = 202553;
```

If the package subprogram changed the row, then ORA_ROWSCN is no longer 9371092, and the update fails.

Your application queries again to obtain the new row data and ORA_ROWSCN. Suppose that the ORA_ROWSCN is now 415639. The application tries the conditional update again, using the new ORA_ROWSCN. This time, the update succeeds, and it is committed. An interactive equivalent is:

```
UPDATE employees
SET salary = salary + 100
WHERE employee_id = 7788
AND ora_rowscn = 415639;
```

ORA_ROWSCN and Tables with Virtual Private Database (VPD)

When a VPD policy is added to a table, it is no longer possible to select the ORA_ROWSCN pseudocolumn. However, because ORA_ROWSCN is available inside the policy function, you can:

1. Create a function that returns a row SCN, as in Example 12–3.

2. In the policy predicate function, add a predicate that stores the row SCN in the context that the function uses while processing rows. For example:

```
||' AND f_ora_rowscn('||object_name||'.ora_rowscn) = 1'
```

3. Use the function to fetch the row. For example:

```
SELECT t.*, get_rowscn(t.rowid) "ORA_ROWSCN" FROM test_table t;
```

Note: To run Example 12–3, you need CREATE ANY CONTEXT system privilege.

Example 12–3 Function that Can Return Row SCN from Table that has VPD

```
-- Create context that function uses while processing rows:

CREATE OR REPLACE FUNCTION f_ora_rowscn
 (rowscn IN NUMBER)
 RETURN NUMBER
AS
BEGIN
  DBMS_SESSION.SET_CONTEXT('STORE_ROWSCN','ROWSCN',rowscn);
  RETURN 1;
END;
/

CREATE CONTEXT store_rowscn USING f_ora_rowscn;

-- Create function that returns row SCN for each row:

CREATE OR REPLACE FUNCTION get_rowscn
  (r IN ROWID)
  RETURN VARCHAR2
AS
BEGIN
  RETURN sys_context('STORE_ROWSCN','ROWSCN');
END;
/
```

Using DBMS_FLASHBACK Package

The DBMS_FLASHBACK package provides the same functionality as Oracle Flashback Query, but Oracle Flashback Query is sometimes more convenient.

The DBMS_FLASHBACK package acts as a time machine: you can turn back the clock, perform normal queries as if you were at that earlier time, and then return to the present. Because you can use the DBMS_FLASHBACK package to perform queries on past data without special clauses such as AS OF or VERSIONS BETWEEN, you can reuse existing PL/SQL code to query the database at earlier times.

You must have the EXECUTE privilege on the DBMS_FLASHBACK package.

To use the DBMS_FLASHBACK package in your PL/SQL code:

1. Specify a past time by invoking either DBMS_FLASHBACK.ENABLE_AT_TIME or DBMS_FLASHBACK.ENABLE_AT_SYSTEM_CHANGE_NUMBER.
2. Perform regular queries (that is, queries without special flashback-feature syntax such as AS OF). Do not perform DDL or DML operations.

The database is queried at the specified past time.

3. Return to the present by invoking DBMS_FLASHBACK.DISABLE.

   You must invoke DBMS_FLASHBACK.DISABLE before invoking DBMS_FLASHBACK.ENABLE_AT_TIME or DBMS_FLASHBACK.ENABLE_AT_SYSTEM_CHANGE_NUMBER again. You cannot nest enable/disable pairs.

To use a cursor to store the results of queries, open the cursor before invoking DBMS_FLASHBACK.DISABLE. After storing the results and invoking DBMS_FLASHBACK.DISABLE, you can:

- Perform INSERT or UPDATE operations to modify the current database state by using the stored results from the past.
- Compare current data with the past data. After invoking DBMS_FLASHBACK.DISABLE, open a second cursor. Fetch from the first cursor to retrieve past data; fetch from the second cursor to retrieve current data. You can store the past data in a temporary table and then use set operators such as MINUS or UNION to contrast or combine the past and current data.

You can invoke DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER at any time to get the current System Change Number (SCN). DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER always returns the current SCN regardless of previous invocations of DBMS_FLASHBACK.ENABLE.

See Also:

- *Oracle Database PL/SQL Packages and Types Reference* for details of the DBMS_FLASHBACK package

Using Flashback Transaction

The DBMS_FLASHBACK.TRANSACTION_BACKOUT procedure rolls back a transaction and its dependent transactions while the database remains online. This recovery operation uses undo data to create and run the **compensating transactions** that return the affected data to its original state.

The transactions being rolled back are subject to these restrictions:

- They cannot have performed DDL operations that changed the logical structure of database tables.
- They cannot use Large Object (LOB) Data Types:
  - BFILE
  - BLOB
  - CLOB
  - NCLOB
- They cannot use features that LogMiner does not support.

  The features that LogMiner supports depends on the value of the COMPATIBLE initialization parameter for the database that is rolling back the transaction. The default value is the release number of the most recent major release.

  Flashback Transaction inherits SQL data type support from LogMiner. Therefore, if LogMiner fails due to an unsupported SQL data type in a the transaction, Flashback Transaction fails too.

Some data types, though supported by LogMiner, do not generate undo information as part of operations that modify columns of such types. Therefore, Flashback Transaction does not support tables containing these data types. These include tables with BLOB, CLOB and XML type.

See Also:

- *Oracle Data Guard Concepts and Administration* for information about data type and DDL support on a logical standby database
- *Oracle Database SQL Language Reference* for information about LOB data types
- *Oracle Database Utilities* for information about LogMiner
- *Oracle Database Administrator's Guide* for information about the `COMPATIBLE` initialization parameter

Topics:

- Dependent Transactions
- TRANSACTION_BACKOUT Parameters
- TRANSACTION_BACKOUT Reports

Dependent Transactions

In the context of Flashback Transaction, transaction 2 can depend on transaction 1 in any of these ways:

- **Write-after-write dependency**

  Transaction 1 changes a row of a table, and later transaction 2 changes the same row.

- **Primary key dependency**

  A table has a primary key constraint on column c. In a row of the table, column c has the value v. Transaction 1 deletes that row, and later transaction 2 inserts a row into the same table, assigning the value v to column c.

- **Foreign key dependency**

  In table b, column b1 has a foreign key constraint on column a1 of table a. Transaction 1 changes a value in a1, and later transaction 2 changes a value in b1.

TRANSACTION_BACKOUT Parameters

The parameters of the `TRANSACTION_BACKOUT` procedure are:

- Number of transactions to be backed out
- List of transactions to be backed out, identified either by name or by XID
- Time hint, if you identify transactions by name

  Specify a time that is earlier than any transaction started.

- Backout option from Table 12–2

For the syntax of the `TRANSACTION_BACKOUT` procedure and detailed parameter descriptions, see *Oracle Database PL/SQL Packages and Types Reference*.

Table 12–2 Flashback TRANSACTION_BACKOUT Options

Option	Description
CASCADE	Backs out specified transactions and all dependent transactions in a post-order fashion (that is, children are backed out before parents are backed out). Without CASCADE, if any dependent transaction is not specified, an error occurs.
NOCASCADE	Default. Backs out specified transactions, which are expected to have no dependent transactions. First dependent transactions causes an error and appears in *_FLASHBACK_TRANSACTION_REPORT.
NOCASCADE_FORCE	Backs out specified transactions, ignoring dependent transactions. Server runs undo SQL statements for specified transactions in reverse order of commit times. If no constraints break and you are satisfied with the result, you can commit the changes; otherwise, you can roll them back.
NONCONFLICT_ONLY	Backs out changes to nonconflicting rows of the specified transactions. Database remains consistent, but transaction atomicity is lost.

TRANSACTION_BACKOUT analyzes the transactional dependencies, performs DML operations, and generates reports. TRANSACTION_BACKOUT does not commit the DML operations that it performs as part of transaction backout, but it holds all the required locks on rows and tables in the right form, preventing other dependencies from entering the system. To make the transaction backout permanent, you must explicitly commit the transaction.

TRANSACTION_BACKOUT Reports

To see the reports that TRANSACTION_BACKOUT generates, query the static data dictionary views *_FLASHBACK_TXN_STATE and *_FLASHBACK_TXN_REPORT.

*_FLASHBACK_TXN_STATE

The static data dictionary view *_FLASHBACK_TXN_STATE shows whether a transaction is active or backed out. If a transaction appears in this view, it is backed out.

*_FLASHBACK_TXN_STATE is maintained atomically for compensating transactions. If a compensating transaction is backed out, all changes that it made are also backed out, and *_FLASHBACK_TXN_STATE reflects this. For example, if compensating transaction ct backs out transactions t1 and t2, then t1 and t2 appear in *_FLASHBACK_TXN_STATE. If ct itself is later backed out, the effects of t1 and t2 are reinstated, and t1 and t2 disappear from *_FLASHBACK_TXN_STATE.

> **See Also:** *Oracle Database Reference* for more information about *_FLASHBACK_TXN_STATE

*_FLASHBACK_TXN_REPORT

The static data dictionary view *_FLASHBACK_TXN_REPORT provides a detailed report for each backed-out transaction.

> **See Also:** *Oracle Database Reference* for more information about *_FLASHBACK_TXN_REPORT

Using Flashback Data Archive (Oracle Total Recall)

A Flashback Data Archive provides the ability to track and store transactional changes to a table over its lifetime. A Flashback Data Archive is useful for compliance with record stage policies and audit reports.

A Flashback Data Archive consists of one or more tablespaces or parts thereof. You can have multiple Flashback Data Archives. If you are logged on as `SYSDBA`, you can specify a default Flashback Data Archive for the system. A Flashback Data Archive is configured with retention time. Data archived in the Flashback Data Archive is retained for the retention time.

By default, flashback archiving is off for any table. You can enable flashback archiving for a table if all of these conditions are true:

- You have the `FLASHBACK ARCHIVE` object privilege on the Flashback Data Archive to use for that table.
- The table is neither nested, clustered, temporary, remote, or external.
- The table contains neither `LONG` nor nested columns.

After flashback archiving is enabled for a table, you can disable it only if you either have the `FLASHBACK ARCHIVE ADMINISTER` system privilege or you are logged on as `SYSDBA`.

When choosing a Flashback Data Archive for a specific table, consider the data retention requirements for the table and the retention times of the Flashback Data Archives on which you have the `FLASHBACK ARCHIVE` object privilege.

Topics:

- Creating a Flashback Data Archive
- Altering a Flashback Data Archive
- Dropping a Flashback Data Archive
- Specifying the Default Flashback Data Archive
- Enabling and Disabling Flashback Data Archive
- DDL Statements on Tables Enabled for Flashback Data Archive
- Viewing Flashback Data Archive Data
- Flashback Data Archive Scenarios

> **See Also:**
>
> `http://www.oracle.com/database/total-recall.html` for more information about Oracle Total Recall

Creating a Flashback Data Archive

Create a Flashback Data Archive with the `CREATE FLASHBACK ARCHIVE` statement, specifying:

- Name of the Flashback Data Archive
- Name of the first tablespace of the Flashback Data Archive
- (Optional) Maximum amount of space that the Flashback Data Archive can use in the first tablespace

The default is unlimited. Unless your space quota on the first tablespace is also unlimited, you must specify this value; otherwise, error ORA-55621 occurs.

- Retention time (number of days that Flashback Data Archive data for the table is guaranteed to be stored)

If you are logged on as SYSDBA, you can also specify that this is the default Flashback Data Archive for the system. If you omit this option, you can still make this Flashback Data Archive the default later (see "Specifying the Default Flashback Data Archive" on page 12-20).

Examples

- Create a default Flashback Data Archive named fla1 that uses up to 10 G of tablespace tbs1, whose data are retained for one year:

```
CREATE FLASHBACK ARCHIVE DEFAULT fla1 TABLESPACE tbs1
  QUOTA 10G RETENTION 1 YEAR;
```

- Create a Flashback Data Archive named fla2 that uses tablespace tbs2, whose data are retained for two years:

```
CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 RETENTION 2 YEAR;
```

For more information about the CREATE FLASHBACK ARCHIVE statement, see *Oracle Database SQL Language Reference*.

Altering a Flashback Data Archive

With the ALTER FLASHBACK ARCHIVE statement, you can:

- Change the retention time of a Flashback Data Archive
- Purge some or all of its data
- Add, modify, and remove tablespaces

Note: Removing all tablespaces of a Flashback Data Archive causes an error.

If you are logged on as SYSDBA, you can also use the ALTER FLASHBACK ARCHIVE statement to make a specific file the default Flashback Data Archive for the system.

Examples

- Make Flashback Data Archive fla1 the default Flashback Data Archive:

```
ALTER FLASHBACK ARCHIVE fla1 SET DEFAULT;
```

- To Flashback Data Archive fla1, add up to 5 G of tablespace tbs3:

```
ALTER FLASHBACK ARCHIVE fla1 ADD TABLESPACE tbs3 QUOTA 5G;
```

- To Flashback Data Archive fla1, add as much of tablespace tbs4 as needed:

```
ALTER FLASHBACK ARCHIVE fla1 ADD TABLESPACE tbs4;
```

- Change the maximum space that Flashback Data Archive fla1 can use in tablespace tbs3 to 20 G:

```
ALTER FLASHBACK ARCHIVE fla1 MODIFY TABLESPACE tbs3 QUOTA 20G;
```

- Allow Flashback Data Archive `fla1` to use as much of tablespace `tbs1` as needed:

  ```
  ALTER FLASHBACK ARCHIVE fla1 MODIFY TABLESPACE tbs1;
  ```

- Change the retention time for Flashback Data Archive `fla1` to two years:

  ```
  ALTER FLASHBACK ARCHIVE fla1 MODIFY RETENTION 2 YEAR;
  ```

- Remove tablespace `tbs2` from Flashback Data Archive `fla1`:

  ```
  ALTER FLASHBACK ARCHIVE fla1 REMOVE TABLESPACE tbs2;
  ```

  (Tablespace `tbs2` is not dropped.)

- Purge all historical data from Flashback Data Archive `fla1`:

  ```
  ALTER FLASHBACK ARCHIVE fla1 PURGE ALL;
  ```

- Purge all historical data older than one day from Flashback Data Archive `fla1`:

  ```
  ALTER FLASHBACK ARCHIVE fla1
    PURGE BEFORE TIMESTAMP (SYSTIMESTAMP - INTERVAL '1' DAY);
  ```

- Purge all historical data older than SCN 728969 from Flashback Data Archive `fla1`:

  ```
  ALTER FLASHBACK ARCHIVE fla1 PURGE BEFORE SCN 728969;
  ```

For more information about the `ALTER FLASHBACK ARCHIVE` statement, see *Oracle Database SQL Language Reference*.

Dropping a Flashback Data Archive

Drop a Flashback Data Archive with the `DROP FLASHBACK ARCHIVE` statement. Dropping a Flashback Data Archive deletes its historical data, but does not drop its tablespaces.

Example

Remove Flashback Data Archive `fla1` and all its historical data, but not its tablespaces:

```
DROP FLASHBACK ARCHIVE fla1;
```

For more information about the `DROP FLASHBACK ARCHIVE` statement, see *Oracle Database SQL Language Reference*.

Specifying the Default Flashback Data Archive

By default, the system has no default Flashback Data Archive. If you are logged on as `SYSDBA`, you can specify default Flashback Data Archive in either of these ways:

- Specify the name of an existing Flashback Data Archive in the `SET DEFAULT` clause of the `ALTER FLASHBACK ARCHIVE` statement. For example:

  ```
  ALTER FLASHBACK ARCHIVE fla1 SET DEFAULT;
  ```

  If `fla1` does not exist, an error occurs.

- Include `DEFAULT` in the `CREATE FLASHBACK ARCHIVE` statement when you create a Flashback Data Archive. For example:

  ```
  CREATE FLASHBACK ARCHIVE DEFAULT fla2 TABLESPACE tbs1
    QUOTA 10G RETENTION 1 YEAR;
  ```

The default Flashback Data Archive for the system is the default Flashback Data Archive for every user who does not have his or her own default Flashback Data Archive.

See Also:

- *Oracle Database SQL Language Reference* for more information about the `CREATE FLASHBACK ARCHIVE` statement
- *Oracle Database SQL Language Reference* for more information about the `ALTER DATABASE` statement

Enabling and Disabling Flashback Data Archive

By default, flashback archiving is disabled for any table. You can enable flashback archiving for a table if you have the `FLASHBACK ARCHIVE` object privilege on the Flashback Data Archive to use for that table.

To enable flashback archiving for a table, include the `FLASHBACK ARCHIVE` clause in either the `CREATE TABLE` or `ALTER TABLE` statement. In the `FLASHBACK ARCHIVE` clause, you can specify the Flashback Data Archive where the historical data for the table are stored. The default is the default Flashback Data Archive for the system. If you specify a nonexistent Flashback Data Archive, an error occurs.

If you enable flashback archiving for a table, but AUM is disabled, error ORA-55614 occurs when you try to modify the table.

If a table has flashback archiving enabled, and you try to enable it again with a different Flashback Data Archive, an error occurs.

After flashback archiving is enabled for a table, you can disable it only if you either have the `FLASHBACK ARCHIVE ADMINISTER` system privilege or you are logged on as `SYSDBA`. To disable flashback archiving for a table, specify `NO FLASHBACK ARCHIVE` in the `ALTER TABLE` statement. (It is unnecessary to specify `NO FLASHBACK ARCHIVE` in the `CREATE TABLE` statement, because that is the default.)

See Also: *Oracle Database SQL Language Reference* for more information about the `FLASHBACK ARCHIVE` clause of the `CREATE TABLE` statement, including restrictions on its use

Examples

- Create table `employee` and store the historical data in the default Flashback Data Archive:

```
CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10),
  JOB VARCHAR2(9), MGR NUMBER(4)) FLASHBACK ARCHIVE;
```

- Create table `employee` and store the historical data in the Flashback Data Archive `fla1`:

```
CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10),
  JOB VARCHAR2(9), MGR NUMBER(4)) FLASHBACK ARCHIVE fla1;
```

- Enable flashback archiving for the table `employee` and store the historical data in the default Flashback Data Archive:

```
ALTER TABLE employee FLASHBACK ARCHIVE;
```

- Enable flashback archiving for the table `employee` and store the historical data in the Flashback Data Archive `fla1`:

  ```
  ALTER TABLE employee FLASHBACK ARCHIVE fla1;
  ```

- Disable flashback archiving for the table `employee`:

  ```
  ALTER TABLE employee NO FLASHBACK ARCHIVE;
  ```

DDL Statements on Tables Enabled for Flashback Data Archive

Flashback Data Archive supports many DDL statements, including some that alter the table definition or move data. For example:

- `ALTER TABLE` statement that does any of the following:
  - Adds, drops, renames, or modifies a column
  - Adds, drops, or renames a constraint
  - Drops or truncates a partition or subpartition operation
- `TRUNCATE TABLE` statement
- `RENAME` statement that renames a table

Some DDL statements cause error ORA-55610 when used on a table enabled for Flashback Data Archive. For example:

- `ALTER TABLE` statement that includes an `UPGRADE TABLE` clause, with or without an `INCLUDING DATA` clause
- `ALTER TABLE` statement that moves or exchanges a partition or subpartition operation
- `DROP TABLE` statement

If you must use unsupported DDL statements on a table enabled for Flashback Data Archive, use the `DBMS_FLASHBACK_ARCHIVE.DISASSOCIATE_FBA` procedure to disassociate the base table from its Flashback Data Archive. To reassociate the Flashback Data Archive with the base table afterward, use the `DBMS_FLASHBACK_ARCHIVE.REASSOCIATE_FBA` procedure.

See Also:

- *Oracle Database SQL Language Reference* for information about the `ALTER TABLE` statement
- *Oracle Database SQL Language Reference* for information about the `TRUNCATE TABLE` statement
- *Oracle Database SQL Language Reference* for information about the `RENAME` statement
- *Oracle Database SQL Language Reference* for information about the `DROP TABLE` statement
- *Oracle Database PL/SQL Packages and Types Reference* for information about the `DBMS_FLASHBACK_ARCHIVE` package

Viewing Flashback Data Archive Data

Table 12–3 lists and briefly describes the static data dictionary views that you can query for information about Flashback Data Archive files.

Table 12–3 Static Data Dictionary Views for Flashback Data Archive Files

View	Description
`*_FLASHBACK_ARCHIVE`	Displays information about Flashback Data Archive files.
`*_FLASHBACK_ARCHIVE_TS`	Displays tablespaces of Flashback Data Archive files.
`*_FLASHBACK_ARCHIVE_TABLES`	Displays information about tables that are enabled for Data Flashback Archive files.

See Also:

- *Oracle Database Reference* for detailed information about `*_FLASHBACK_ARCHIVE`
- *Oracle Database Reference* for detailed information about `*_FLASHBACK_ARCHIVE_TS`
- *Oracle Database Reference* for detailed information about `*_FLASHBACK_ARCHIVE_TABLES`

Flashback Data Archive Scenarios

- Scenario: Using Flashback Data Archive to Enforce Digital Shredding
- Scenario: Using Flashback Data Archive to Access Historical Data
- Scenario: Using Flashback Data Archive to Generate Reports
- Scenario: Using Flashback Data Archive for Auditing
- Scenario: Using Flashback Data Archive to Recover Data

Scenario: Using Flashback Data Archive to Enforce Digital Shredding

Your company wants to "shred" (delete) historical data changes to the `Taxes` table after ten years. When you create the Flashback Data Archive for `Taxes`, you specify a retention time of ten years:

```
CREATE FLASHBACK ARCHIVE taxes_archive TABLESPACE tbs1 RETENTION 10 YEAR;
```

When history data from transactions on `Taxes` exceeds the age of ten years, it is purged. (The `Taxes` table itself, and history data from transactions less than ten years old, are not purged.)

Scenario: Using Flashback Data Archive to Access Historical Data

You want to be able to retrieve the inventory of all items at the beginning of the year from the table `inventory`, and to be able to retrieve the stock price for each symbol in your portfolio at the close of business on any specified day of the year from the table `stock_data`.

Create a default Flashback Data Archive named `fla1` that uses up to 10 G of tablespace `tbs1`, whose data are retained for five years (you must be logged on as `SYSDBA`):

```
CREATE FLASHBACK ARCHIVE DEFAULT fla1 TABLESPACE tbs1
  QUOTA 10G RETENTION 5 YEAR;
```

Enable Flashback Data Archive for the tables `inventory` and `stock_data`, and store the historical data in the default Flashback Data Archive:

```
ALTER TABLE inventory FLASHBACK ARCHIVE;
```

```
ALTER TABLE stock_data FLASHBACK ARCHIVE;
```

To retrieve the inventory of all items at the beginning of the year 2007, use this query:

```
SELECT product_number, product_name, count FROM inventory AS OF
  TIMESTAMP TO_TIMESTAMP ('2007-01-01 00:00:00', 'YYYY-MM-DD HH24:MI:SS');
```

To retrieve the stock price for each symbol in your portfolio at the close of business on July 23, 2007, use this query:

```
SELECT symbol, stock_price FROM stock_data AS OF
  TIMESTAMP TO_TIMESTAMP ('2007-07-23 16:00:00', 'YYYY-MM-DD HH24:MI:SS')
  WHERE symbol IN my_portfolio;
```

Scenario: Using Flashback Data Archive to Generate Reports

You want users to be able to generate reports from the table `investments`, for data stored in the past five years.

Create a default Flashback Data Archive named `fla2` that uses up to 20 G of tablespace `tbs1`, whose data are retained for five years (you must be logged on as `SYSDBA`):

```
CREATE FLASHBACK ARCHIVE DEFAULT fla2 TABLESPACE tbs1
  QUOTA 20G RETENTION 5 YEAR;
```

Enable Flashback Data Archive for the table `investments`, and store the historical data in the default Flashback Data Archive:

```
ALTER TABLE investments FLASHBACK ARCHIVE;
```

Lisa wants a report on the performance of her investments at the close of business on December 31, 2006. She uses this query:

```
SELECT * FROM investments AS OF
  TIMESTAMP TO_TIMESTAMP ('2006-12-31 16:00:00', 'YYYY-MM-DD HH24:MI:SS')
  WHERE name = 'LISA';
```

Scenario: Using Flashback Data Archive for Auditing

A medical insurance company must audit a medical clinic. The medical insurance company has its claims in the table `Billings`, and creates a default Flashback Data Archive named `fla4` that uses up to 100 G of tablespace `tbs1`, whose data are retained for 10 years:

```
CREATE FLASHBACK ARCHIVE DEFAULT fla4 TABLESPACE tbs1
  QUOTA 100G RETENTION 10 YEAR;
```

The company enables Flashback Data Archive for the table `Billings`, and stores the historical data in the default Flashback Data Archive:

```
ALTER TABLE Billings FLASHBACK ARCHIVE;
```

On May 1, 2007, clients were charged the wrong amounts for some diagnoses and tests. To see the records as of May 1, 2007, the company uses this query:

```
SELECT date_billed, amount_billed, patient_name, claim_Id,
  test_costs, diagnosis FROM Billings AS OF TIMESTAMP
  TO_TIMESTAMP('2007-05-01 00:00:00', 'YYYY-MM-DD HH24:MI:SS');
```

Scenario: Using Flashback Data Archive to Recover Data

An end user recovers from erroneous transactions that were previously committed in the database. The undo data for the erroneous transactions is no longer available, but because the required historical information is available in the Flashback Data Archive, Flashback Query works seamlessly.

Lisa manages a software development group whose product sales are doing well. On November 3, 2007, she decides to give all her level-three employees who have more than two years of experience a salary increase of 10% and a promotion to level four. Lisa asks her HR representative, Bob, to make the changes.

Using the HR web application, Bob updates the `employee` table to give Lisa's level-three employees a 10% raise and a promotion to level four. Then Bob finishes his work for the day and leaves for home, unaware that he omitted the requirement of two years of experience in his transaction. A few days later, Lisa checks to see if Bob has done the updates and finds that everyone in the group was given a raise! She calls Bob immediately and asks him to correct the error.

At first, Bob thinks he cannot return the employee table to its prior state without going to the backups. Then he remembers that the `employee` table has Flashback Data Archive enabled.

First, he verifies that no other transaction modified the `employee` table after his: The commit time stamp from the transaction query corresponds to Bob's transaction, two days ago.

Next, Bob uses these statements to return the `employee` table to the way it was before his erroneous change:

```
DELETE EMPLOYEE WHERE MANAGER = 'LISA JOHNSON';
INSERT INTO EMPLOYEE
  SELECT * FROM EMPLOYEE
    AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '2' DAY)
      WHERE MANAGER = 'LISA JOHNSON';
```

Bob then reexecutes the update that Lisa had requested.

General Guidelines for Oracle Flashback Technology

- Use the `DBMS_FLASHBACK.ENABLE` and `DBMS_FLASHBACK.DISABLE` procedures around SQL code that you do not control, or when you want to use the same past time for several consecutive queries.
- Use Oracle Flashback Query, Oracle Flashback Version Query, or Oracle Flashback Transaction Query for SQL code that you write, for convenience. An Oracle Flashback Query, for example, is flexible enough to do comparisons and store results in a single query.
- To obtain an SCN to use later with a flashback feature, use `DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER`.
- To compute or retrieve a past time to use in a query, use a function return value as a time stamp or SCN argument. For example, add or subtract an `INTERVAL` value to the value of the `SYSTIMESTAMP` function.
- Use Oracle Flashback Query, Oracle Flashback Version Query, and Oracle Flashback Transaction Query locally or remotely. An example of a remote Oracle Flashback Query is:

  ```
  (SELECT * FROM employees@some_remote_host AS OF
      TIMESTAMP (SYSTIMESTAMP - INTERVAL '60' MINUTE);
  ```

- To ensure database consistency, always perform a `COMMIT` or `ROLLBACK` operation before querying past data.
- Remember that all flashback processing uses the current session settings, such as national language and character set, not the settings that were in effect at the time being queried.
- Remember that DDLs that alter the structure of a table (such as drop/modify column, move table, drop partition, truncate table/partition, and add constraint) invalidate any existing undo data for the table. If you try to retrieve data from a time before such a DDL executed, error ORA-01466 occurs. DDL operations that alter the storage attributes of a table (such as `PCTFREE`, `INITRANS`, and `MAXTRANS`) do not invalidate undo data.
- To query past data at a precise time, use an SCN. If you use a time stamp, the actual time queried might be up to 3 seconds earlier than the time you specify. Oracle Database uses SCNs internally and maps them to time stamps at a granularity of 3 seconds.

  For example, suppose that the SCN values 1000 and 1005 are mapped to the time stamps 8:41 AM and 8:46 AM, respectively. A query for a time between 8:41:00 and 8:45:59 AM is mapped to SCN 1000; an Oracle Flashback Query for 8:46 AM is mapped to SCN 1005. Therefore, if you specify a time that is slightly after a DDL operation (such as a table creation) Oracle Database might use an SCN that is just before the DDL operation, causing error ORA-01466.
- You cannot retrieve past data from a dynamic performance (`V$`) view. A query on such a view always returns current data.
- You can perform queries on past data in static data dictionary views, such as `*_TABLES`.

Performance Guidelines for Oracle Flashback Technology

- Use the `DBMS_STATS` package to generate statistics for all tables involved in an Oracle Flashback Query. Keep the statistics current. Oracle Flashback Query uses the cost-based optimizer, which relies on these statistics.
- Minimize the amount of undo data that must be accessed. Use queries to select small sets of past data using indexes, not to scan entire tables. If you must scan a full table, add a parallel hint to the query.

  The performance cost in I/O is the cost of paging in data and undo blocks that are not in the buffer cache. The performance cost in CPU use is the cost of applying undo information to affected data blocks. When operating on changes in the recent past, flashback operations are CPU-bound.
- For Oracle Flashback Version Query, use index structures. Oracle Database keeps undo data for index changes and data changes. Performance of index lookup-based Oracle Flashback Version Query is an order of magnitude faster than the full table scans that are otherwise needed.
- In an Oracle Flashback Transaction Query, the `xid` column is of the type `RAW(8)`. To take advantage of the index built on the `xid` column, use the `HEXTORAW` conversion function: `HEXTORAW(xid)`.
- A Oracle Flashback Query against a materialized view does not take advantage of query rewrite optimization.

13

Choosing a Programming Environment

To choose a programming environment for a development project, read:

- The topics in this chapter and the documents to which they refer.
- The platform-specific documents that explain which compilers and development tools your platforms support.

Sometimes the choice of programming environment is obvious, for example:

- Pro*COBOL does not support ADTs or collection types, while Pro*C/C++ does.
- SQLJ does not support dynamic SQL the way that JDBC does.

If no programming language provides all the features you need, you can use multiple programming languages, because:

- Every programming language in this chapter can invoke PL/SQL and Java stored subprograms. (Stored subprograms include triggers and ADT methods.)
- PL/SQL, Java, SQL, and Oracle Call Interface (OCI) can invoke external C subprograms.
- External C subprograms can access Oracle Database using SQL, OCI, or Pro*C (but not C++).

For more information about multilanguage programming, see Chapter 14, "Developing Applications with Multiple Programming Languages."

Topics:

- Overview of Application Architecture
- Overview of the Program Interface
- Overview of PL/SQL
- Overview of Oracle Database Java Support
- Choosing PL/SQL or Java
- Overview of Precompilers
- Overview of OCI and OCCI
- Choosing a Precompiler or OCI
- Overview of Oracle Data Provider for .NET (ODP.NET)
- Overview of OraOLEDB
- Overview of Oracle Objects for OLE (OO4O)

Overview of Application Architecture

In this topic, **application architecture** refers to the computing environment in which a database application connects to an Oracle Database.

Topics:

- Client/Server Architecture
- Server-Side Programming
- Two-Tier and Three-Tier Architecture

> **See Also:** *Oracle Database Concepts* for more information about application architecture

Client/Server Architecture

In a traditional client/server program, your application code runs on a client system; that is, a system other than the database server. Database calls are transmitted from the client system to the database server. Data is transmitted from the client to the server for insert and update operations and returned from the server to the client for query operations. The data is processed on the client system. Client/server programs are typically written by using precompilers, whereas SQL statements are embedded within the code of another language such as C, C++, or COBOL.

> **See Also:** *Oracle Database Concepts* for more information about client/server architecture

Server-Side Programming

You can develop application logic that resides entirely inside the database by using triggers that are executed automatically when changes occur in the database or stored subprograms that are invoked explicitly. Off-loading the work from your application lets you reuse code that performs verification and cleanup and control database operations from a variety of clients. For example, by making stored subprograms invocable through a Web server, you can construct a Web-based user interface that performs the same functions as a client/server application.

> **See Also:** *Oracle Database Concepts* for more information about server-side programming

Two-Tier and Three-Tier Architecture

Client/server computing is often referred to as a **two-tier model**: your application communicates directly with the database server. In the **three-tier model**, a separate application server processes the requests. The application server might be a basic Web server, or might perform advanced functions like caching and load-balancing. Increasing the processing power of this middle tier lets you lessen the resources needed by client systems, resulting in a **thin client configuration** in which the client system might need only a Web browser or other means of sending requests over the TCP/IP or HTTP protocols.

> **See Also:** *Oracle Database Concepts* for more information about multitier architecture

Overview of the Program Interface

The **program interface** is the software layer between a database application and Oracle Database. The program interface:

- Provides a security barrier, preventing destructive access to the SGA by client user processes
- Acts as a communication mechanism, formatting information requests, passing data, and trapping and returning errors
- Converts and translates data, particularly between different types of computers or to external user program data types

The Oracle code acts as a server, performing database tasks on behalf of an application (a client), such as fetching rows from data blocks. The program interface consists of several parts, provided by both Oracle Database software and operating system-specific software.

> **See Also:** *Oracle Database Concepts* for more information about the program interface

Topics:

- User Interface
- Stateful and Stateless User Interfaces

User Interface

The **user interface** is what your application displays to end users. It depends on the technology behind the application and the needs of the users themselves. Experienced users can enter SQL statements that are passed on to the database. Novice users can be shown a graphical user interface that uses the graphics libraries of the client system (such as Windows or X-Windows). Any of these traditional user interfaces can also be provided in a Web browser through HTML and Java.

Stateful and Stateless User Interfaces

In traditional client/server applications, the application can keep a record of user actions and use this information over the course of one or more sessions. For example, past choices can be presented in a menu so that they do not have to be entered again. When the application can save information in this way, the application is considered **stateful**.

Web or thin-client applications that are **stateless** are easier to develop. Stateless applications gather all the required information, process it using the database, and then start over with the next user. This is a popular way to process single-screen requests such as customer registration.

There are many ways to add stateful action to Web applications that are stateless by default. For example, an entry form on one Web page can pass information to subsequent Web pages, enabling you to construct a wizard-like interface that remembers user choices through several different steps. You can use cookies to store small items of information about the client system, and retrieve them when the user returns to a Web site. You can use servlets to keep a database session open and store variables between requests from the same client.

Overview of PL/SQL

PL/SQL, the Oracle procedural extension of SQL, is a completely portable, high-performance transaction-processing language. PL/SQL lets you manipulate data with SQL statements; control program flow with conditional selection and loops; declare constants and variables; define subprograms; define types, subtypes, and ADTs and declare variables of those types; and trap run-time errors.

Applications written in any of the Oracle Database programmatic interfaces can invoke PL/SQL stored subprograms and send blocks of PL/SQL code to Oracle Database for execution. Third-generation language (3GL) applications can access PL/SQL scalar and composite data types through host variables and implicit data type conversion. A 3GL language is easier than assembler language for a human to understand and includes features such as named variables. Unlike a fourth-generation language (4GL), it is not specific to an application domain.

You can use PL/SQL to develop stored procedures that can be invoked by a Web client.

See Also:

- *Oracle Database PL/SQL Language Reference* for information about the advantages, main features, and architecture of PL/SQL
- Chapter 9, "Developing PL/SQL Web Applications," to learn how to use PL/SQL in Web development

Overview of Oracle Database Java Support

This section provides an overview of built-in database features that support Java applications. The database includes the core JDK libraries such as `java.lang`, `java.io`, and so on. The database supports client-side Java standards such as JDBC and SQLJ, and provides server-side JDBC and SQLJ drivers that enable data-intensive Java code to run within the database.

Topics:

- Overview of Oracle JVM
- Overview of Oracle JDBC
- Overview of Oracle SQLJ
- Comparing Oracle JDBC and Oracle SQLJ
- Overview of Oracle JPublisher
- Overview of Java Stored Subprograms
- Overview of Oracle Database Web Services

See Also:

- *Oracle Database Java Developer's Guide*
- *Oracle Database JDBC Developer's Guide*
- *Oracle Database JPublisher User's Guide*

Overview of Oracle JVM

Oracle JVM, the Java Virtual Machine provided with the Oracle Database, is compliant with the J2SE version 1.5.x specification and supports the database session architecture.

Any database session can activate a dedicated JVM. All sessions share the same JVM code and statics; however, private states for any given session are held, and subsequently garbage collected, in an individual session space.

This design provides these benefits:

- Java applications have the same session isolation and data integrity as SQL operations.
- You need not run Java in a separate process for data integrity.
- Oracle JVM is a robust JVM with a small memory footprint.
- The JVM has the same linear Symmetric Multiprocessing (SMP) scalability as the database and can support thousands of concurrent Java sessions.

Oracle JVM works consistently with every platform supported by Oracle Database. Java applications that you develop with Oracle JVM can easily be ported to any supported platform.

Oracle JVM includes a deployment-time native compiler that enables Java code to be compiled once, stored in executable form, shared among users, and invoked more quickly and efficiently.

Security features of the database are also available with Oracle JVM. Java classes must be loaded in a database schema (by using Oracle JDeveloper, a third-party IDE, SQL*Plus, or the `loadjava` utility) before they can be called. Java class calls are secured and controlled through database authentication and authorization, Java 2 security, and invoker's rights (IR) or definer's rights (DR).

> **See Also:** *Oracle Database Concepts* for additional general information about Oracle JVM

Overview of Oracle JDBC

Java Database Connectivity (JDBC) is an Applications Programming Interface (API) that enables Java to send SQL statements to an object-relational database such as Oracle Database.

Oracle Database includes these extensions to the JDBC 1.22 standard:

- Support for Oracle data types
- Performance enhancement by row prefetching
- Performance enhancement by execution batching
- Specification of query column types to save round-trips
- Control of `DatabaseMetaData` calls

Oracle Database supports all APIs from the JDBC 2.0 standard, including the core APIs, optional packages, and numerous extensions. Some highlights include datasources, JTA, and distributed transactions.

Oracle Database supports these features from the JDBC 3.0 standard:

- Support for JDK 1.5.
- Toggling between local and global transactions.

- Transaction savepoints.
- Reuse of prepared statements by connection pools.

> **Note:** JDBC code and SQLJ code interoperate. For more information, see "Comparing Oracle JDBC and Oracle SQLJ" on page 13-10.)

Topics:

- Oracle JDBC Drivers
- Sample JDBC 2.0 Program
- Sample Pre-2.0 JDBC Program

> **See Also:** *Oracle Database Concepts* for additional general information about Java support in Oracle Database

Oracle JDBC Drivers

The JDBC standard defines four types of JDBC drivers:

Type	Description
1	A JDBC-ODBC bridge. Software must be installed on client systems.
2	Native methods (calls C or C++) and Java methods. Software must be installed on the client.
3	Pure Java. The client uses sockets to call middleware on the server.
4	The most pure Java solution. Talks directly to the database by using Java sockets.

JDBC is based on Part 3 of the SQL standard, "Call-Level Interface."

You can use JDBC to do dynamic SQL. In dynamic SQL, the embedded SQL statement to be executed is not known before the application is run and requires input to build the statement.

The drivers that are implemented by Oracle have extensions to the capabilities in the JDBC standard that was defined by Sun Microsystems.

Topics:

- JDBC Thin Driver
- JDBC OCI Driver
- JDBC Server-Side Internal Driver

> **See Also:**
>
> - *Oracle Database Concepts* for additional general information about JDBC drivers
> - *Oracle Database JDBC Developer's Guide* for more information about JDBC

JDBC Thin Driver The JDBC thin driver is a Type 4 (100% pure Java) driver that uses Java sockets to connect directly to a database server. It has its own implementation of a Two-Task Common (TTC), a lightweight implementation of TCP/IP from Oracle Net. It is written entirely in Java and is therefore platform-independent.

The thin driver does not require Oracle software on the client side. It does need a TCP/IP listener on the server side. Use this driver in Java applets that are downloaded into a Web browser or in applications for which you do not want to install Oracle client software. The thin driver is self-contained, but it opens a Java socket, and thus can only run in a browser that supports sockets.

JDBC OCI Driver The JDBC OCI driver is a Type 2 JDBC driver. It makes calls to OCI written in C to interact with Oracle Database, thus using native and Java methods.

The OCI driver provides access to more features than the thin driver, such as Transparent Application Fail-Over, advanced security, and advanced LOB manipulation.

The OCI driver provides the highest compatibility between different Oracle Database versions. It also supports all installed Oracle Net adapters, including IPC, named pipes, TCP/IP, and IPX/SPX.

Because it uses native methods (a combination of Java and C) the OCI driver is platform-specific. It requires a client installation of version Oracle8*i* or later including Oracle Net, OCI libraries, CORE libraries, and all other dependent files. The OCI driver usually runs faster than the thin driver.

The OCI driver is not appropriate for Java applets, because it uses a C library that is platform-specific and cannot be downloaded into a Web browser. It is usable in J2EE components running in middle-tier application servers, such as Oracle Application Server. Oracle Application Server provides middleware services and tools that support access between applications and browsers.

JDBC Server-Side Internal Driver The JDBC server-side internal driver is a Type 2 driver that runs inside the database server, reducing the number of round-trips needed to access large amounts of data. The driver, the Java server VM, the database, the Java native compiler (which speeds execution by as much as 10 times), and the SQL engine all run within the same address space.

This driver provides server-side support for any Java program used in the database. You can also call PL/SQL stored subprograms and triggers.

The server driver fully supports the same features and extensions as the client-side drivers.

Sample JDBC 2.0 Program

This example shows the recommended technique for looking up a data source using JNDI in JDBC 2.0:

```
// import the JDBC packages
import java.sql.*;
import javax.sql.*;
import oracle.jdbc.pool.*;
...
   InitialContext ictx = new InitialContext();
   DataSource ds = (DataSource)ictx.lookup("jdbc/OracleDS");
   Connection conn = ds.getConnection();
   Statement stmt = conn.createStatement();
   ResultSet rs = stmt.executeQuery("SELECT last_name FROM employees");
   while ( rs.next() ) {
   out.println( rs.getString("ename") + "<br>");
   }
conn.close();
```

Sample Pre-2.0 JDBC Program

This source code registers an Oracle JDBC thin driver, connects to the database, creates a `Statement` object, runs a query, and processes the result set.

The `SELECT` statement retrieves and lists the contents of the `last_name` column of the `hr.employees` table.

```
import java.sql.*
import java.math.*
import java.io.*
import java.awt.*

class JdbcTest {
  public static void main (String args []) throws SQLException {
    // Load Oracle driver
    DriverManager.registerDriver (new oracle.jdbc.OracleDriver());

    // Connect to the local database
    Connection conn =
      DriverManager.getConnection ("jdbc:oracle:thin:@myhost:1521:orcl",
                                   "hr", "password");

    // Query the employee names
    Statement stmt = conn.createStatement ();
    ResultSet rset = stmt.executeQuery ("SELECT last_name FROM employees");

    // Print the name out
    while (rset.next ())
      System.out.println (rset.getString (1));
    // Close the result set, statement, and the connection
    rset.close();
    stmt.close();
    conn.close();
  }
}
```

One Oracle Database extension to the JDBC drivers is a form of the `getConnection()` method that uses a `Properties` object. The `Properties` object lets you specify user, password, database information, row prefetching, and execution batching.

To use the OCI driver in this code, replace the `Connection` statement with this code, where `MyHostString` is an entry in the `tnsnames.ora` file:

```
Connection conn = DriverManager.getConnection ("jdbc:oracle:oci8:@MyHostString",
    "hr", "password");
```

If you are creating an applet, then the `getConnection()` and `registerDriver()` strings are different.

Overview of Oracle SQLJ

> **Note:** In this document, **SQLJ** refers to Oracle SQLJ and its extensions.

SQLJ is an ANSI SQL-1999 standard for embedding SQL statements in Java source code. SQLJ provides a simpler alternative to JDBC for both client-side and server-side SQL data access from Java.

A SQLJ source file contains Java source with embedded SQL statements. Oracle SQLJ supports dynamic and static SQL. Support for dynamic SQL is an Oracle extension to the SQLJ standard.

Oracle Database provides a translator and a run time driver to support SQLJ. The SQLJ translator is 100% pure Java and is portable to any JVM that is compliant with JDK version 1.1 or higher.

The Oracle SQLJ translator performs these tasks:

- Translates SQLJ source to Java code with calls to the SQLJ run time driver. The SQLJ translator converts the source code to pure Java source code and can check the syntax and semantics of static SQL statements against a database schema and verify the type compatibility of host variables with SQL types.
- Compiles the generated Java code with the Java compiler.
- (Optional) Creates profiles for the target database. SQLJ generates "profile" files with customization specific to Oracle Database.

Oracle Database supports SQLJ stored subprograms and triggers that run in the Oracle JVM. SQLJ is integrated with JDeveloper. Source-level debugging support for SQLJ is available in JDeveloper.

This is an example of a simple SQLJ executable statement, which returns one value because `employee_id` is unique in the `employee` table:

```
String name;
#sql  { SELECT first_name INTO :name FROM employees WHERE employee_id=112 };
System.out.println("Name is " + name + ", employee number = " + employee_id);
```

Each host variable (or qualified name or complex Java host expression) included in a SQL expression is preceded by a colon (:). Other SQLJ statements declare Java types. For example, you can declare an iterator (a construct related to a database cursor) for queries that retrieve many values, as follows:

```
#sql iterator EmpIter (String EmpNam, int EmpNumb);
```

> **See Also:** *Oracle Database JPublisher User's Guide* for more examples and details about Oracle SQLJ syntax

Topics:

- Benefits of SQLJ
- SQLJ Stored Subprograms in the Server

> **See Also:** *Oracle Database Concepts* for additional general information about SQLJ

Benefits of SQLJ

Oracle SQLJ extensions to Java enable rapid development and easy maintenance of applications that perform database operations through embedded SQL.

In particular, Oracle SQLJ does this:

- Provides a concise, legible mechanism for database access from static SQL. Most SQL in applications is static. SQLJ provides more concise and less error-prone static SQL constructs than JDBC does.
- Provides an SQL Checker module for verification of syntax and semantics at translate time.

- Provides flexible deployment configurations, which makes it possible to implement SQLJ on the client, server, or middle tier.
- Supports a software standard. SQLJ is an effort of a group of vendors and is supported by all of them. Applications can access multiple database vendors.
- Provides source code portability. Executables can be used with all of the vendor DBMSs if the code does not rely on any vendor-specific features.
- Enforces a uniform programming style for the clients and the servers.
- Integrates the SQLJ translator with **Oracle JDeveloper**, a graphical IDE that provides SQLJ translation, Java compilation, profile customizing, and debugging at the source code level, all in one step.
- Includes Oracle Database type extensions.

SQLJ Stored Subprograms in the Server

SQLJ applications can be stored and executed in the server by using these techniques:

- Translate, compile, and customize the SQLJ source code on a client and load the generated classes and resources into the server with the `loadjava` utility. The classes are typically stored in a Java archive (`.jar`) file.
- Load the SQLJ source code into the server, also using `loadjava`, where it is translated and compiled by the server's embedded translator.

Comparing Oracle JDBC and Oracle SQLJ

JDBC code and SQLJ code interoperate, enabling dynamic SQL statements in JDBC to be used with both static and dynamic SQL statements in SQLJ. A SQLJ iterator class corresponds to the JDBC result set.

Some differences between JDBC and SQLJ are:

- JDBC provides a complete dynamic SQL interface from Java to databases. It gives developers full control over database operations. SQLJ simplifies Java database programming to improve development productivity.
- JDBC provides fine-grained control of the execution of dynamic SQL from Java, whereas SQLJ provides a higher-level binding to SQL operations in a specific database schema.
- SQLJ source code is more concise than equivalent JDBC source code.
- SQLJ uses database connections to type-check static SQL code. JDBC, being a completely dynamic API, does not.
- SQLJ provides strong typing of query outputs and return parameters and provides type-checking on calls. JDBC passes values to and from SQL without compile-time type checking.
- SQLJ programs enable direct embedding of Java bind expressions within SQL statements. JDBC requires a separate get or set statement for each bind variable and specifies the binding by position number.
- SQLJ provides simplified rules for calling SQL stored subprograms.

  For example, the following four examples show, on successive lines, how to call a stored procedure or a stored function using either JDBC escape syntax or Oracle JDBC syntax:

  ```
  prepStmt.prepareCall("{call fun(?,?)}");       //stored proc. JDBC esc.
  ```

```
prepStmt.prepareCall("{? = call fun(?,?)}");   //stored func. JDBC esc.
prepStmt.prepareCall("begin fun(:1,:2);end;"); //stored proc. Oracle
prepStmt.prepareCall("begin :1 := fun(:2,:3);end;"); //stored func. Oracle
```

The SQLJ equivalent is:

```
#sql {call fun(param_list) };  //Stored procedure
// Declare x
...
#sql x = {VALUES(fun(param_list)) };  // Stored function
// where VALUES is the SQL construct
```

These benefits are common to SQLJ and JDBC:

- SQLJ source files can contain JDBC calls. SQLJ and JDBC are interoperable.
- Oracle JPublisher generates custom Java classes to be used in your SQLJ or JDBC application for mappings to Oracle Database ADTs and collections.
- PL/SQL and Java stored subprograms can be used interchangeably.

Overview of Oracle JPublisher

Oracle JPublisher is a code generator that automates the process of creating database-centric Java classes by hand. Oracle JPublisher is a client-side utility and is built into the database system. You can run Oracle JPublisher from the command line or directly from the Oracle JDeveloper IDE.

Oracle JPublisher inspects PL/SQL packages and database object types such as ADTs, `VARRAY` types, and nested table types, and then generates a Java class that is a wrapper around the PL/SQL package with corresponding fields and methods.

The generated Java class can be incorporated and used by Java clients or J2EE components to exchange and transfer database object type instances to and from the database transparently.

See Also:

- *Oracle Database Concepts* for additional general information about Oracle JPublisher
- *Oracle Database JPublisher User's Guide* for complete information about Oracle JPublisher

Overview of Java Stored Subprograms

Java stored subprograms enable you to implement programs that run in the database server and are independent of programs that run in the middle tier. Structuring applications in this way reduces complexity and increases reuse, security, performance, and scalability.

For example, you can create a Java stored subprogram that performs operations that require data persistence and a separate program to perform presentation or business logic operations.

Java stored subprograms interface with SQL using an execution model similar to that of PL/SQL.

See Also:

- *Oracle Database Concepts* for additional general information about Java stored subprograms
- *Oracle Database Java Developer's Guide* for complete information about Java stored subprograms

Overview of Oracle Database Web Services

Web services represent a distributed computing paradigm for Java application development that is an alternative to earlier Java protocols such as JDBC, and which enable applications to interact through the XML and Web protocols. For example, an electronics parts vendor can provide a Web-based programmatic interface to its suppliers for inventory management. The vendor can invoke a Web service as part of a program and automatically order stock based on the data returned.

The key technologies used in Web services are:

- Web Services Description Language (WSDL), which is a standard format for creating an XML document. WSDL describes what a web service can do, where it resides, and how to invoke it. Specifically, it describes the operations and parameters, including parameter types, provided by a Web service. In addition, a WSDL document describes the location, the transport protocol, and the invocation style for the Web service.
- Simple Object Access Protocol (SOAP) messaging, which is an XML-based message protocol used by Web services. SOAP does not prescribe a specific transport mechanism such as HTTP, FTP, SMTP, or JMS; however, most Web services accept messages that use HTTP or HTTPS.
- Universal Description, Discovery, and Integration (UDDI) business registry, which is a directory that lists Web services on the internet. The UDDI registry is often compared to a telephone directory, listing unique identifiers (white pages), business categories (yellow pages), and instructions for binding to a service protocol (green pages).

Web services can use a variety of techniques and protocols. For example:

- Dispatching can occur in a synchronous (typical) or asynchronous manner.
- You can invoke a Web service in an RPC-style operation in which arguments are sent and a response returned, or in a message style such as a one-way SOAP document exchange.
- You can use different encoding rules: literal or encoded.

You can invoke a Web service statically, when you might know everything about it beforehand, or dynamically, in which case you can discover its operations and transport endpoints while using it.

Oracle Database can function as either a Web service provider or as a Web service consumer. When used as a provider, the database enables sharing and disconnected access to stored subprograms, data, metadata, and other database resources such as the queuing and messaging systems.

As a Web service provider, Oracle Database provides a disconnected and heterogeneous environment that:

- Exposes stored subprograms independently of the language in which the subprograms are written
- Exposes SQL Queries and XQuery

See Also: *Oracle Database Concepts* for additional general information about Oracle Database as a Web service provider

Choosing PL/SQL or Java

PL/SQL and Java interoperate in the server. You can run a PL/SQL package from Java or wrap a PL/SQL class with a Java wrapper so that it can be invoked from distributed CORBA and Enterprise Java Beans clients.

Table 13–1 shows PL/SQL packages and their Java equivalents.

Table 13–1 PL/SQL Packages and Their Java Equivalents

PL/SQL Package	Java Equivalent
DBMS_ALERT	Call package with SQLJ or JDBC.
DBMS_DDL	JDBC has this functionality.
DBMS_JOB	Schedule a job that has a Java stored subprogram.
DBMS_LOCK	Call with SQLJ or JDBC.
DBMS_MAIL	Use JavaMail.
DBMS_OUTPUT	Use subclass oracle.aurora.rdbms.OracleDBMSOutputStream or Java stored subprogram DBMS_JAVA.SET_STREAMS.
DBMS_PIPE	Call with SQLJ or JDBC.
DBMS_SESSION	Use JDBC to run an ALTER SESSION statement.
DBMS_SNAPSHOT	Call with SQLJ or JDBC.
DBMS_SQL	Use JDBC.
DBMS_TRANSACTION	Use JDBC to run an ALTER SESSION statement.
DBMS_UTILITY	Call with SQLJ or JDBC.
UTL_FILE	Grant the JAVAUSERPRIV privilege and then use Java I/O entry points.

Topics:

- Similarities of PL/SQL and Java
- PL/SQL Advantages Over Java
- Java Advantages Over PL/SQL

Similarities of PL/SQL and Java

Both PL/SQL and Java have built-in packages and libraries.

Both PL/SQL and Java have object-oriented features:

- Both have inheritance.
- PL/SQL has **type evolution**, the ability to change methods and attributes of a type while preserving subtypes and table data that use the type.
- Java has polymorphism and component models for developing distributed systems.

PL/SQL Advantages Over Java

As an extension of SQL, PL/SQL supports all SQL data types, data encapsulation, information hiding, overloading, and exception-handling. Therefore:

- SQL data types are easier to use in PL/SQL than in Java.
- SQL operations are faster with PL/SQL than with Java, especially when a large amount of data is involved, when mostly database access is done, or when bulk operations are used.

  Some advanced PL/SQL capabilities are not available for Java in Oracle9*i* (for example, autonomous transactions and the dblink facility for remote databases).

Code development is usually faster in PL/SQL than in Java.

Java Advantages Over PL/SQL

Java is used for open distributed applications, and many Java-based development tools are available throughout the industry. Java has a richer type system than PL/SQL. Java can use CORBA (which can have many different computer languages in its clients) and Enterprise Java Beans. PL/SQL packages can be invoked from CORBA or Enterprise Java Beans clients. You can run XML tools, the Internet File System, or JavaMail from Java.

Overview of Precompilers

Client/server programs are typically written using **precompilers**, which are programming tools that let you embed SQL statements in high-level programs written in languages such as C, C++, or COBOL. Because the client application hosts the SQL statements, it is called a **host program**, and the language in which it is written is called the **host language**.

A precompiler accepts the host program as input, translates the embedded SQL statements into standard database run-time library calls, and generates a source program that you can compile, link, and run in the usual way.

Topics:

- Overview of the Pro*C/C++ Precompiler
- Overview of the Pro*COBOL Precompiler

> **See Also:** *Oracle Database Concepts* for additional general information about Oracle precompilers

Overview of the Pro*C/C++ Precompiler

For the Pro*C/C++ precompiler, the host language is either C or C++. Some features of the Pro*C/C++ precompiler are:

- You can write multithreaded programs if your platform supports a threads package. Concurrent connections are supported in either single-threaded or multithreaded applications.
- You can improve performance by embedding PL/SQL blocks. These blocks can invoke subprograms in Java or PL/SQL that are written by you or provided in Oracle Database packages.
- Using precompiler options, you can check the syntax and semantics of your SQL or PL/SQL statements during precompilation, and at run time.

- You can invoke stored PL/SQL and Java subprograms. Modules written in COBOL or in C can be invoked from Pro*C/C++. External C subprograms in shared libraries can be invoked by your program.
- You can conditionally precompile sections of your code so that they can run in different environments.
- You can use arrays, or structures, or arrays of structures as host and indicator variables in your code to improve performance.
- You can deal with errors and warnings so that data integrity is guaranteed. As a programmer, you control how errors are handled.
- Your program can convert between internal data types and C language data types.
- The Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI), lower-level C and C++ interfaces, are available for use in your precompiler source.
- Pro*C/C++ supports dynamic SQL, a technique that enables users to input variable values and statement syntax.
- Pro*C/C++ can use special SQL statements to manipulate tables containing user-defined object types. An Object Type Translator (OTT) maps the ADTs and named collection types in your database to structures and headers that you include in your source.
- Two kinds of collection types, nested tables and `VARRAY`, are supported with a set of SQL statements that give you a high degree of control over data.
- Large Objects are accessed by another set of SQL statements.
- A new ANSI SQL standard for dynamic SQL is supported for new applications, so that you can run SQL statements with a varying number of host variables. An older technique for dynamic SQL is still usable by pre-existing applications.
- Globalization support lets you use multibyte characters and UCS2 Unicode data.
- Using scrollable cursors, you can move backward and forward through a result set. For example, you can fetch the last row of the result set, or jump forward or backward to an absolute or relative position within the result set.
- A **connection pool** is a group of physical connections to a database that can be shared by several named connections. Enabling the connection pool option can help optimize the performance of Pro*C/C++ application. The connection pool option is not enabled by default.

> **See Also:** *Pro*C/C++ Precompiler Programmer's Guide* for complete information about the Pro*C/C++ precompiler

Example 13–1 is a code fragment from a C source program that queries the table `employees` in the schema `hr`.

Example 13–1 Pro*C/C++ Application

```
...
#define  UNAME_LEN   10
...
int   emp_number;
/* Define a host structure for the output values of a SELECT statement. */
/* No declare section needed if precompiler option MODE=ORACLE          */
struct {
    VARCHAR  last_name[UNAME_LEN];
    float    salary;
```

```
    float    commission_pct;
} emprec;
/* Define an indicator structure to correspond to the host output structure. */
struct {
    short emp_name_ind;
    short sal_ind;
    short comm_ind;
} emprec_ind;
...
/* Select columns last_name, salary, and commission_pct given the user's input
/* for employee_id. */
    EXEC SQL SELECT last_name, salary, commission_pct
        INTO :emprec INDICATOR :emprec_ind
        FROM employees
        WHERE employee_id = :emp_number;
...
```

The embedded `SELECT` statement differs slightly from the interactive (SQL*Plus) `SELECT` statement. Every embedded SQL statement begins with `EXEC SQL`. The colon (:) precedes every host (C) variable. The returned values of data and indicators (set when the data value is `NULL` or character columns were truncated) can be stored in structs (such as in the preceding code fragment), in arrays, or in arrays of structs. Multiple result set values are handled very simply in a manner that resembles the case shown, where there is only one result, because of the unique employee number. Use the actual names of columns and tables in embedded SQL.

Either use the default precompiler option values or enter values that give you control over the use of resources, how errors are reported, the formatting of output, and how cursors (which correspond to a particular connection or SQL statement) are managed. Cursors are used when there are multiple result set values.

Enter the options either in a configuration file, on the command line, or inline inside your source code with a special statement that begins with `EXEC ORACLE`. If there are no errors found, you can compile, link, and run the output source file, like any other C program that you write.

Use the precompiler to create server database access from clients that can be on many different platforms. Pro*C/C++ gives you the freedom to design your own user interfaces and to add database access to existing applications.

Before writing your embedded SQL statements, you can test interactive versions of the SQL in SQL*Plus and then make minor changes to start testing your embedded SQL application.

Overview of the Pro*COBOL Precompiler

For the Pro*COBOL precompiler, the host language is COBOL. Some features of the Pro*COBOL precompiler are:

- You can invoke stored PL/SQL or Java subprograms. You can improve performance by embedding PL/SQL blocks. These blocks can invoke PL/SQL subprograms written by you or provided in Oracle Database packages.
- Precompiler options enable you to define how cursors, errors, syntax-checking, file formats, and so on, are handled.
- Using precompiler options, you can check the syntax and semantics of your SQL or PL/SQL statements during precompilation, and at run time.
- You can conditionally precompile sections of your code so that they can run in different environments.

- Use tables, or group items, or tables of group items as host and indicator variables in your code to improve performance.
- You can program how errors and warnings are handled, so that data integrity is guaranteed.
- Pro*COBOL supports dynamic SQL, a technique that enables users to input variable values and statement syntax.

> **See Also:** *Pro*COBOL Programmer's Guide* for complete information about the Pro*COBOL precompiler

Example 13–2 is a code fragment from a COBOL source program that queries the table `employees` in the schema `hr`.

Example 13–2 Pro*COBOL Application

```
...
 WORKING-STORAGE SECTION.
*
* DEFINE HOST INPUT AND OUTPUT HOST AND INDICATOR VARIABLES.
* NO DECLARE SECTION NEEDED IF MODE=ORACLE.
*
 01  EMP-REC-VARS.
     05  EMP-NAME    PIC X(10) VARYING.
     05  EMP-NUMBER  PIC S9(4) COMP VALUE ZERO.
     05  SALARY      PIC S9(5)V99 COMP-3 VALUE ZERO.
     05  COMMISSION  PIC S9(5)V99 COMP-3 VALUE ZERO.
     05  COMM-IND    PIC S9(4) COMP VALUE ZERO.
...
 PROCEDURE DIVISION.
...
     EXEC SQL
         SELECT last_name, salary, commission_pct
         INTO :EMP-NAME, :SALARY, :COMMISSION:COMM-IND
         FROM employees
         WHERE employee_id = :EMP-NUMBER
     END-EXEC.
...
```

The embedded `SELECT` statement is only slightly different from an interactive (SQL*Plus) `SELECT` statement. Every embedded SQL statement begins with EXEC SQL. The colon (:) precedes every host (COBOL) variable. The SQL statement is terminated by `END-EXEC`. The returned values of data and indicators (set when the data value is `NULL` or character columns were truncated) can be stored in group items (such as in the preceding code fragment), in tables, or in tables of group items. Multiple result set values are handled very simply in a manner that resembles the case shown, where there is only one result, given the unique employee number. Use the actual names of columns and tables in embedded SQL.

Use the default precompiler option values, or enter values that give you control over the use of resources, how errors are reported, the formatting of output, and how cursors are managed (cursors correspond to a particular connection or SQL statement).

Enter the options in a configuration file, on the command line, or inline inside your source code with a special statement that begins with EXEC ORACLE. If there are no errors found, you can compile, link, and run the output source file, like any other COBOL program that you write.

Use the precompiler to create server database access from clients that can be on many different platforms. Pro*COBOL gives you the freedom to design your own user interfaces and to add database access to existing COBOL applications.

The embedded SQL statements available conform to an ANSI standard, so that you can access data from many databases in a program, including remote servers networked through Oracle Net.

Before writing your embedded SQL statements, you can test interactive versions of the SQL in SQL*Plus and then make minor changes to start testing your embedded SQL application.

Overview of OCI and OCCI

The Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI) are application programming interfaces (APIs) that enable you to create applications that use native subprogram invocations of a third-generation language to access Oracle Database and control all phases of SQL statement execution. These APIs provide:

- Improved performance and scalability through the efficient use of system memory and network connectivity
- Consistent interfaces for dynamic session and transaction management in a two-tier client/server or multitier environment
- *N*-tiered authentication
- Comprehensive support for application development using Oracle Database objects
- Access to external databases
- Ability to develop applications that service an increasing number of users and requests without additional hardware investments

OCI lets you manipulate data and schemas in a database using a host programming language, such as C. OCCI is an object-oriented interface suitable for use with C++. These APIs provide a library of standard database access and retrieval functions in the form of a dynamic run-time library (OCILIB) that can be linked in an application at run time. You need not embed SQL or PL/SQL within 3GL programs.

See Also: For more information about OCI and OCCI calls:

- *Oracle Call Interface Programmer's Guide*
- *Oracle C++ Call Interface Programmer's Guide*
- *Oracle Streams Advanced Queuing User's Guide*
- *Oracle Database Globalization Support Guide*
- *Oracle Database Data Cartridge Developer's Guide*

Topics:

- Advantages of OCI and OCCI
- OCI and OCCI Functions
- Procedural and Nonprocedural Elements of OCI and OCCI Applications
- Building an OCI or OCCI Application

Advantages of OCI and OCCI

OCI and OCCI provide significant advantages over other methods of accessing Oracle Database:

- More fine-grained control over all aspects of the application design.
- High degree of control over program execution.
- Use of familiar 3GL programming techniques and application development tools such as browsers and debuggers.
- Support of dynamic SQL, method 4.
- Availability on the broadest range of platforms of all the Oracle Database programmatic interfaces.
- Dynamic bind and define using callbacks.
- Describe functionality to expose layers of server metadata.
- Asynchronous event notification for registered client applications.
- Enhanced array data manipulation language (DML) capability for arrays.
- Ability to associate a commit request with an run to reduce round-trips.
- Optimization for queries using transparent prefetch buffers to reduce round-trips.
- Thread safety, so you do not have to implement mutual exclusion (mutex) locks on OCI and OCCI handles.
- The server connection in nonblocking mode means that control returns to the OCI or OCCI code when a call is still running or cannot complete.

OCI and OCCI Functions

Both OCI and OCCI have four kinds of functions:

Kind of Function	Purpose
Relational	To manage database access and process SQL statements
Navigational	To manipulate objects retrieved from the database
Database mapping and manipulation	To manipulate data attributes of Oracle Database types
External subprogram	To write C callbacks from PL/SQL

Procedural and Nonprocedural Elements of OCI and OCCI Applications

OCI and OCCI enable you to develop applications that combine the nonprocedural data access power of SQL with the procedural capabilities of most programming languages, including C and C++. Procedural and nonprocedural languages have these characteristics:

- In a nonprocedural language program, the set of data to be operated on is specified, but what operations are performed and how the operations are to be carried out is not specified. The nonprocedural nature of SQL makes it an easy language to learn and to use to perform database transactions. It is also the standard language used to access and manipulate data in modern relational and object-relational database systems.

- In a procedural language program, the execution of most statements depends on previous or subsequent statements and on control structures, such as loops or conditional branches, which are not available in SQL. The procedural nature of these languages makes them more complex than SQL, but it also makes them very flexible and powerful.

The combination of both nonprocedural and procedural language elements in an OCI or OCCI program provides easy access to Oracle Database in a structured programming environment.

OCI and OCCI support all SQL data definition, data manipulation, query, and transaction control facilities that are available through Oracle Database. For example, an OCI or OCCI program can run a query against Oracle Database. The queries can require the program to supply data to the database using input (bind) variables, as follows:

```
SELECT name FROM employees WHERE empno = :empnumber
```

In the preceding SQL statement, `:empnumber` is a placeholder for a value to be supplied by the application.

Alternatively, you can use PL/SQL, Oracle's procedural extension to SQL. The applications you develop can be more powerful and flexible than applications written in SQL alone. OCI and OCCI also provide facilities for accessing and manipulating objects in Oracle Database.

Building an OCI or OCCI Application

As Figure 13–1 shows, you compile and link an OCI or OCCI program in the same way that you compile and link a nondatabase application. There is no need for a separate preprocessing or precompilation step.

Figure 13–1 The OCI or OCCI Development Process

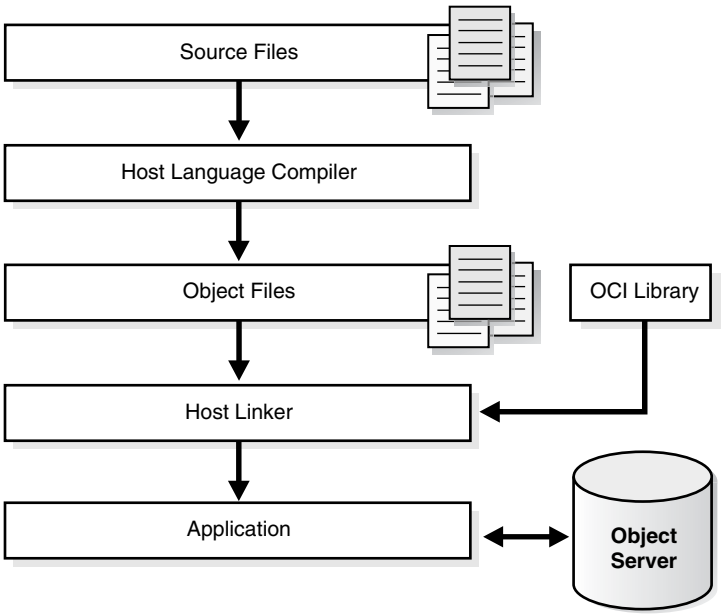

Note: To properly link your OCI and OCCI programs, it might be necessary on some platforms to include other libraries, in addition to the OCI and OCCI libraries. Check your Oracle platform-specific documentation for further information about extra libraries that might be required.

Choosing a Precompiler or OCI

Precompiler applications typically contain less code than equivalent OCI applications, which can help productivity.

Some situations require detailed control of the database and are suited for OCI applications (either pure OCI or a precompiler application with embedded OCI calls):

- OCI provides more detailed control over multiplexing and migrating sessions.
- OCI provides dynamic bind and define using callbacks that can be used for any arbitrary structure, including lists.
- OCI has many calls to handle metadata.
- OCI enables asynchronous event notifications to be received by a client application. It provides a means for clients to generate notifications for propagation to other clients.
- OCI enables DML statements to use arrays to complete as many iterations as possible before returning any error messages.
- OCI calls for special purposes include Advanced Queuing, globalization support, Data Cartridges, and support of the date and time data types.
- OCI calls can be embedded in a Pro*C/C++ application.

Overview of Oracle Data Provider for .NET (ODP.NET)

Oracle Data Provider for .NET (ODP.NET) is an implementation of a data provider for Oracle Database.

ODP.NET uses APIs native to Oracle Database to offer fast and reliable access from any .NET application to database features and data. It also uses and inherits classes and interfaces available in the Microsoft .NET Framework Class Library.

For programmers using Oracle Provider for OLE DB, ADO (ActiveX Data Objects) provides an automation layer that exposes an easy programming model. ADO.NET provides a similar programming model, but without the automation layer, for better performance. More importantly, the ADO.NET model enables native providers such as ODP.NET to expose specific features and data types specific to Oracle Database.

See Also: *Oracle Data Provider for .NET Developer's Guide*

This is a simple C# application that connects to Oracle Database and displays its version number before disconnecting:

```
using System;
using Oracle.DataAccess.Client;

class Example
{
  OracleConnection con;
```

```
  void Connect()
  {
    con = new OracleConnection();
    con.ConnectionString = "User Id=hr;Password=password;Data Source=oracle";
    con.Open();
    Console.WriteLine("Connected to Oracle" + con.ServerVersion);
  }

  void Close()
  {
    con.Close();
    con.Dispose();
  }

  static void Main()
  {
    Example example = new Example();
    example.Connect();
    example.Close();
  }
}
```

Note: Additional samples are provided in directory *ORACLE_BASE**ORACLE_HOME*\ODP.NET\Samples.

Overview of OraOLEDB

Oracle Provider for OLE DB (OraOLEDB) is an OLE DB data provider that offers high performance and efficient access to Oracle data by OLE DB consumers. In general, this developer's guide assumes that you are using OraOLEDB through OLE DB or ADO.

See Also: *Oracle Provider for OLE DB Developer's Guide for Microsoft Windows*

Overview of Oracle Objects for OLE (OO4O)

Oracle Objects for OLE (OO4O) is a product designed to provide easy access to data stored in Oracle Database with any programming or scripting language that supports the Microsoft COM Automation and ActiveX technology. This includes Visual Basic, Visual C++, Visual Basic For Applications (VBA), IIS Active Server Pages (VBScript and JavaScript), and others.

See the OO4O online help for detailed information about using OO4O.

Oracle Objects for OLE consists of these software layers:

- OO4O "In-Process" Automation Server
- Oracle Data Control
- Oracle Objects for OLE C++ Class Library

Figure 13–2 illustrates the OO4O software components.

Figure 13–2 Software Layers

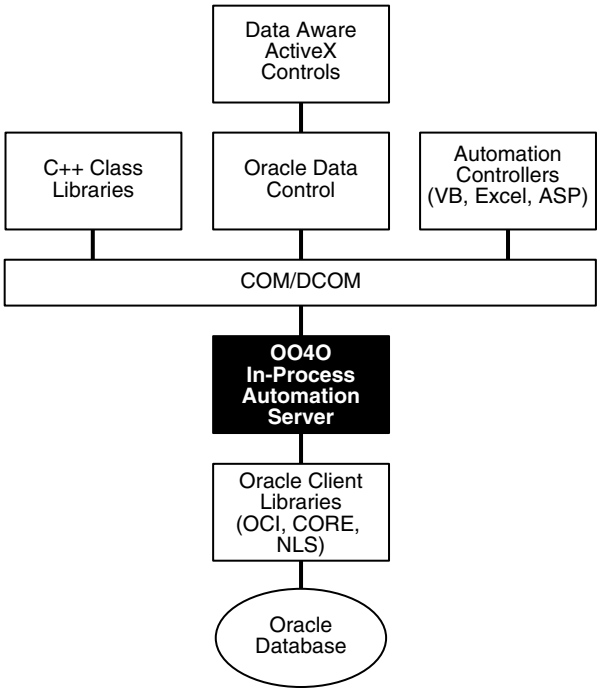

Topics:

- OO4O Automation Server
- OO4O Object Model
- Support for Oracle LOB and Object Data Types
- Oracle Data Control
- Oracle Objects for OLE C++ Class Library

OO4O Automation Server

The OO4O Automation Server is a set of COM Automation objects for connecting to Oracle Database, running SQL statements and PL/SQL blocks, and accessing the results.

Unlike other COM-based database connectivity APIs, such as Microsoft ADO, the OO4O Automation Server was developed specifically for use with Oracle Database.

It provides an optimized API for accessing features that are unique to Oracle Database and are otherwise cumbersome or inefficient to use from ODBC or OLE database-specific components.

OO4O provides key features for accessing Oracle Database efficiently and easily in environments ranging from the typical two-tier client/server applications, such as those developed in Visual Basic or Excel, to application servers deployed in multitiered application server environments such as Web server applications in Microsoft Internet Information Server (IIS) or Microsoft Transaction Server (MTS).

Features include:

- Support for execution of PL/SQL and Java stored subprograms, and PL/SQL anonymous blocks. This includes support for Oracle Database data types used as

parameters to stored subprograms, including PL/SQL cursors. See "Support for Oracle LOB and Object Data Types" on page 13-28.

- Support for scrollable and updatable cursors for easy and efficient access to result sets of queries.
- Thread-safe objects and Connection Pool Management Facility for developing efficient Web server applications.
- Full support for Oracle Database object-relational and LOB data types.
- Full support for Advanced Queuing.
- Support for array inserts and updates.
- Support for Microsoft Transaction Server (MTS).

OO4O Object Model

The Oracle Objects for OLE object model is illustrated in Figure 13–3.

Figure 13–3 Objects and Their Relations

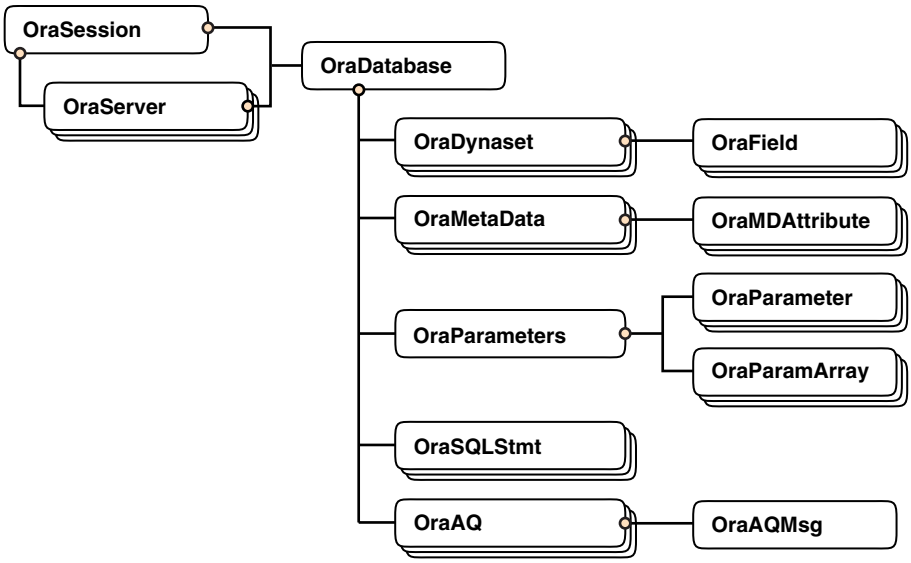

Topics:

- OraSession
- OraServer
- OraDatabase
- OraDynaset
- OraField
- OraMetaData and OraMDAttribute
- OraParameter and OraParameters
- OraParamArray
- OraSQLStmt
- OraAQ

- OraAQMsg
- OraAQAgent

OraSession

An OraSession object manages collections of OraDatabase, OraConnection, and OraDynaset objects used within an application.

Typically, a single OraSession object is created for each application, but you can create named OraSession objects for shared use within and between applications.

The OraSession object is the top-most object for an application. It is the only object created by the CreateObject VB/VBA API and not by an Oracle Objects for OLE method. This code fragment shows how to create an OraSession object:

```
Dim OraSession as Object
Set OraSession = CreateObject("OracleInProcServer.XOraSession")
```

OraServer

OraServer represents a physical network connection to Oracle Database.

The `OraServer` interface is introduced to expose the connection-multiplexing feature provided in the Oracle Call Interface. After an `OraServer` object is created, multiple user sessions (`OraDatabase`) can be attached to it by calling the `OpenDatabase` method. This feature is particularly useful for application components, such as Internet Information Server (IIS), that use Oracle Objects for OLE in n-tier distributed environments.

The use of connection multiplexing when accessing Oracle Database with a large number of user sessions active can help reduce server processing and resource requirements while improving server scalability.

OraServer is used to share a single connection across multiple OraDatabase objects (multiplexing), whereas each OraDatabase obtained from an OraSession has its own physical connection.

OraDatabase

An `OraDatabase` interface adds additional methods for controlling transactions and creating interfaces representing Oracle Database object types. Attributes of schema objects can be retrieved using the `Describe` method of the `OraDatabase` interface.

Before Oracle Database 8*g* Release 1 (8.1), an `OraDatabase` object is created by calling the `OpenDatabase` method of an `OraSession` interface. The Oracle Net alias, user name, and password are passed as arguments to this method. As of Oracle Database 8*g* Release 1 (8.1), calling this method results in implicit creation of an `OraServer` object.

An `OraDatabase` object can also be created using the `OpenDatabase` method of the `OraServer` interface.

Transaction control methods are available at the `OraDatabase` (user session) level. Transactions might be started as `Read-Write` (default), `Serializable`, or `Read-only`. Transaction control methods include:

- `BeginTrans`
- `CommitTrans`
- `RollbackTrans`

For example:

```
UserSession.BeginTrans(OO4O_TXN_READ_WRITE)
UserSession.ExecuteSQL("delete emp where empno = 1234")
UserSession.CommitTrans
```

OraDynaset

An `OraDynaset` object permits browsing and updating of data created from a SQL `SELECT` statement.

The `OraDynaset` object can be thought of as a cursor, although in actuality several real cursors might be used to implement the semantics of `OraDynaset`. An `OraDynaset` object automatically maintains a local cache of data fetched from the server and transparently implements scrollable cursors within the browse data. Large queries might require significant local disk space; application developers are encouraged to refine queries to limit disk usage.

OraField

An `OraField` object represents a single column or data item within a row of a dynaset.

If the current row is being updated, then the `OraField` object represents the currently updated value, although the value might not have been committed to the database.

Assignment to the `Value` property of a field is permitted only if a record is being edited (using `Edit`) or a record is being added (using `AddNew`). Other attempts to assign data to a field's `Value` property results in an error.

OraMetaData and OraMDAttribute

An `OraMetaData` object is a collection of `OraMDAttribute` objects that represent the description information about a particular schema object in the database.

The `OraMetaData` object can be visualized as a table with three columns:

- `Metadata Attribute Name`
- `Metadata Attribute Value`
- `Flag` specifying whether the `Value` is another `OraMetaData` object

The `OraMDAttribute` objects contained in the `OraMetaData` object can be accessed by subscripting using ordinal integers or by using the name of the property. Referencing a subscript that is not in the collection results in the return of a `NULL` `OraMDAttribute` object.

OraParameter and OraParameters

An `OraParameter` object represents a bind variable in a SQL statement or PL/SQL block.

`OraParameter` objects are created, accessed, and removed indirectly through the `OraParameters` collection of an `OraDatabase` object. Each parameter has an identifying name and an associated value. You can automatically bind a parameter to SQL and PL/SQL statements of other objects (as noted in the object descriptions), by using the parameter name as a placeholder in the SQL or PL/SQL statement. Such use of parameters can simplify dynamic queries and increase program performance.

OraParamArray

An `OraParamArray` object represents an array-type bind variable in a SQL statement or PL/SQL block, as opposed to a scalar-type bind variable represented by the `OraParameter` object.

`OraParamArray` objects are created, accessed, and removed indirectly through the `OraParameters` collection of an `OraDatabase` object. Each `OraParamArray` object has an identifying name and an associated value.

OraSQLStmt

An `OraSQLStmt` object represents a single SQL statement. Use the `CreateSQL` method to create an `OraSQLStmt` object from an `OraDatabase` object.

During create and refresh, `OraSQLStmt` objects automatically bind all relevant, enabled input parameters to the specified SQL statement, using the parameter names as placeholders in the SQL statement. This can improve the performance of SQL statement execution without reparsing the SQL statement.

The `OraSQLStmt` object can be used later to run the same query using a different value for the :`SALARY` placeholder. This is done as follows (`updateStmt` is the `OraSQLStmt` object here):

```
OraDatabase.Parameters("SALARY").value = 200000
updateStmt.Parameters("ENAME").value = "KING"
updateStmt.Refresh
```

OraAQ

An `OraAQ` object is instantiated by calling the `CreateAQ` method of the `OraDatabase` interface. It represents a queue that is present in the database.

Oracle Objects for OLE provides interfaces for accessing Oracle Advanced Queuing (AQ) feature. It makes AQ accessible from popular COM-based development environments such as Visual Basic. For a detailed description of Oracle Advanced Queuing, see *Oracle Streams Advanced Queuing User's Guide*.

OraAQMsg

The `OraAQMsg` object encapsulates the message to be enqueued or dequeued. The message can be of any user-defined or raw type.

For a detailed description of Oracle Advanced Queuing, see *Oracle Streams Advanced Queuing User's Guide.*

OraAQAgent

The `OraAQAgent` object represents a message recipient and is only valid for queues that support multiple consumers. It is a child of `OraAQMsg`.

An `OraAQAgent` object can be instantiated by calling the `AQAgent` method. For example:

```
Set agent = qMsg.AQAgent(name)
```

An `OraAQAgent` object can also be instantiated by calling the `AddRecipient` method. For example:

```
Set agent = qMsg.AddRecipient(name, address, protocol).
```

Support for Oracle LOB and Object Data Types

Oracle Objects for OLE (OO4O) provides full support for accessing and manipulating instances of object data types and LOBs in Oracle Database. Figure 13–4 illustrates the data types supported by OO4O.

Instances of these types can be fetched from the database or passed as input or output variables to SQL statements and PL/SQL blocks, including stored subprograms. All instances are mapped to COM Automation Interfaces that provide methods for dynamic attribute access and manipulation.

Figure 13–4 Supported Oracle Database Data Types

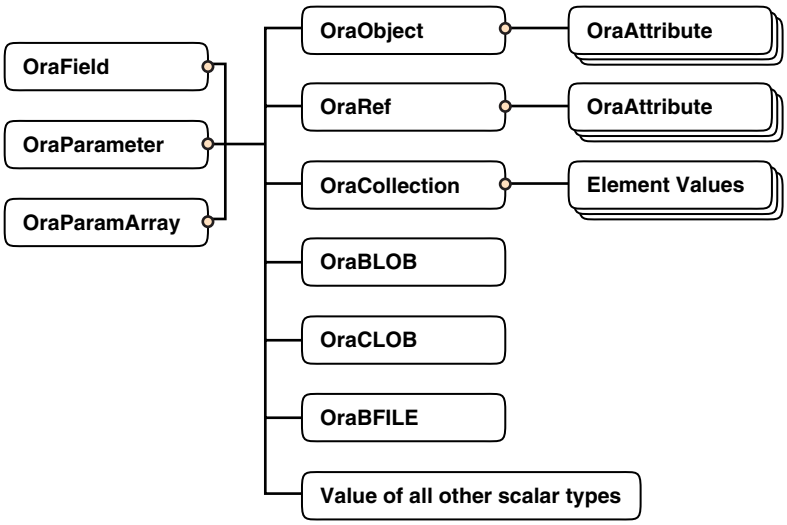

Topics:

- OraBLOB and OraCLOB
- OraBFILE

OraBLOB and OraCLOB

The `OraBlob` and `OraClob` interfaces in Oracle Objects for OLE provide methods for performing operations on large database objects of data type `BLOB`, `CLOB`, and `NCLOB`. `BLOB`, `CLOB`, and `NCLOB` data types are also referred to here as **LOB** data types.

LOB data is accessed using `Read` and the `CopyToFile` methods.

LOB data is modified using `Write`, `Append`, `Erase`, `Trim`, `Copy`, `CopyFromFile`, and `CopyFromBFile` methods. Before modifying the content of a LOB column in a row, a row lock must be obtained. If the LOB column is a field of an `OraDynaset`, object, then the lock is obtained by calling the `Edit` method.

OraBFILE

The `OraBFile` interface in Oracle Objects for OLE provides methods for performing operations on large database objects of data type `BFILE`.

`BFILE` objects are large binary data objects stored in operating system files outside of the database tablespaces.

Oracle Data Control

Oracle Data Control (ODC) is an ActiveX Control that is designed to simplify the exchange of data between Oracle Database and visual controls such edit, text, list, and grid controls in Visual Basic and other development tools that support custom controls.

ODC acts as an agent to handle the flow of information from Oracle Database and a visual data-aware control, such as a grid control, that is bound to it. The data control manages various user interface (UI) tasks such as displaying and editing data. It also runs and manages the results of database queries.

Oracle Data Control is compatible with the Microsoft data control included with Visual Basic. If you are familiar with the Visual Basic data control, learning to use Oracle Data Control is quick and easy. Communication between data-aware controls and a Data Control is governed by a protocol that Microsoft specified.

Oracle Objects for OLE C++ Class Library

Oracle Objects for OLE (OO4O) C++ Class Library is a collection of C++ classes that provide programmatic access to the Oracle Object Server. Although the class library is implemented using OLE Automation, neither the OLE development kit nor any OLE development knowledge is necessary to use it. This library helps C++ developers avoid the chore of writing COM client code for accessing the OO4O interfaces.

See Also:

- For detailed information about Oracle Objects for OLE see the online help provided with the OO4O product:

  Oracle Objects for OLE Help

  Oracle Objects for OLE C++ Class Library Help

- For examples of how to use Oracle Objects for OLE:

  Samples in the `ORACLE_HOME\OO4O` directory of the Oracle Database installation

  Oracle Database SecureFiles and Large Objects Developer's Guide

  Oracle Streams Advanced Queuing User's Guide

14

Developing Applications with Multiple Programming Languages

This chapter explains how you can develop database applications that call external procedures written in other programming languages.

Topics:

- Overview of Multilanguage Programs
- What Is an External Procedure?
- Overview of Call Specification for External Procedures
- Loading External Procedures
- Publishing External Procedures
- Publishing Java Class Methods
- Publishing External C Procedures
- Locations of Call Specifications
- Passing Parameters to External C Procedures with Call Specifications
- Running External Procedures with CALL Statements
- Handling Errors and Exceptions in Multilanguage Programs
- Using Service Routines with External C Procedures
- Doing Callbacks with External C Procedures

Overview of Multilanguage Programs

Oracle Database lets you work in different languages:

- PL/SQL, as described in the *Oracle Database PL/SQL Language Reference*
- C, through the Oracle Call Interface (OCI), as described in the *Oracle Call Interface Programmer's Guide*
- C or C++, through the Pro*C/C++ precompiler, as described in the *Pro*C/C++ Programmer's Guide*
- COBOL, through the Pro*COBOL precompiler, as described in the *Pro*COBOL Programmer's Guide*
- .NET , through Oracle Data Provider for .NET (ODP.NET), as described in *Oracle Data Provider for .NET Developer's Guide*

- Visual Basic, through Oracle Objects for OLE (OO4O)and Oracle Provider for OLE DB, as described in *Oracle Objects for OLE Developer's Guide for Microsoft Windows* and *Oracle Provider for OLE DB Developer's Guide for Microsoft Windows*.
- Java, through the JDBC Application Programmers Interface (API). See *Oracle Database Java Developer's Guide*.

How can you choose between these different implementation possibilities? Each of these languages offers different advantages: ease of use, the availability of programmers with specific expertise, the need for portability, and the existence of legacy code are powerful determinants.

The choice might narrow depending on how your application must work with Oracle Database:

- PL/SQL is a powerful development tool, specialized for SQL transaction processing.
- Some computation-intensive tasks are executed most efficiently in a lower level language, such as C.
- For both portability and security, you might select Java.

Most significantly for performance, only PL/SQL and Java methods run within the address space of the server. C/C++ methods are dispatched as external procedures, and run on the server system but outside the address space of the database server. Pro*COBOL and Pro*C/C++ are precompilers, and Visual Basic accesses Oracle Database through the OCI, which is implemented in C.

Taking all these factors into account suggests that there might be situations in which you might need to implement your application in multiple languages. For example, the introduction of Java running within the address space of the server suggest that you might want to import existing Java applications into the database, and then leverage this technology by calling Java functions from PL/SQL and SQL.

PL/SQL external procedures enable you to write C procedure calls as PL/SQL bodies. These C procedures are callable directly from PL/SQL, and from SQL through PL/SQL procedure calls. The database provides a special-purpose interface, the call specification, that lets you call external procedures from other languages. While this service is designed for intercommunication between SQL, PL/SQL, C, and Java, it is accessible from any base language that can call these languages. For example, your procedure can be written in a language other than Java or C, and if C can call your procedure, then SQL or PL/SQL can use it. Therefore, if you have a candidate C++ procedure, use a C++ `extern "C"` statement in that procedure to make it callable by C.

Therefore, the strengths and capabilities of different languages are available to you, regardless of your programmatic environment. You are not restricted to one language with its inherent limitations. External procedures promote reusability and modularity because you can deploy specific languages for specific purposes.

What Is an External Procedure?

An **external procedure** is a procedure stored in a dynamic link library (DLL), or libunit for a Java class method. You register the procedure with the base language, and then call it to perform special-purpose processing.

For example, when you work in PL/SQL, the language loads the library dynamically at run time, and then calls the procedure as if it were a PL/SQL procedure. These

procedures participate fully in the current transaction and can call back to the database to perform SQL operations.

The procedures are loaded only when necessary, so memory is conserved. Because the decoupling of the call specification from its implementation body means that the procedures can be enhanced without affecting the calling programs.

External procedures let you:

- Isolate execution of client applications and processes from the database instance to ensure that any problems on the client side do not adversely impact the database.
- Move computation-bound programs from client to server where they run faster (because they avoid the round-trips of network communication)
- Interface the database server with external systems and data sources
- Extend the functionality of the database server itself

> **Note:** The external library (DLL file) must be statically linked. In other words, it must not reference any external symbols from other external libraries (DLL files). Oracle Database does not resolve such symbols, so they can cause your external procedure to fail.

Overview of Call Specification for External Procedures

You publish external procedures through **call specifications**, which provide a superset of the `AS EXTERNAL` function through the `AS LANGUAGE` clause. `AS LANGUAGE` call specifications allow the publishing of external C procedures, but also Java class methods.

> **Note:** To support legacy applications, call specifications also enable you to publish with the `AS EXTERNAL` clause. For application development, however, using the `AS LANGUAGE` clause is recommended.

In general, call specifications enable:

- Dispatching the appropriate C or Java target procedure
- Data type conversions
- Parameter mode mappings
- Automatic memory allocation and cleanup
- Purity constraints to be specified, where necessary, for package functions called from SQL.
- Calling Java methods or C procedures from database triggers
- Location flexibility: you can put `AS LANGUAGE` call specifications in package or type specifications, or package (or type) bodies to optimize performance and hide implementation details

To use an existing program as an external procedure, load, publish, and then call it.

Loading External Procedures

To make your external C procedures or Java methods available to PL/SQL, you must first load them. The manner of doing this depends upon whether the procedure is written in C or Java.

Topics:

- Loading Java Class Methods
- Loading External C Procedures

Loading Java Class Methods

One way to load Java programs is to use the `CREATE JAVA` statement, which you can run interactively from SQL*Plus. When implicitly called by the `CREATE JAVA` statement, the Java Virtual Machine (JVM)] library manager loads Java binaries (`.class` files) and resources from local `BFILE` or LOB columns into RDBMS libunits.

Suppose a compiled Java class is stored in the operating system file `/home/java/bin/Agent.class`.

Create a class `libunit` in schema *`username`* from file `Agent.class` as follows:

1. Connect to the database as `SYSTEM` and grant the user *`username`* the `CREATE ANY DIRECTORY` privilege.
2. Connect to the database as *`username`* and create a directory object on the server's file system:

   ```
   CREATE DIRECTORY Bfile_dir AS '/home/java/bin';
   ```

   The name of the directory object is an alias for the directory path leading to `Agent.class`.
3. Create the class `libunit`:

   ```
   CREATE JAVA CLASS USING BFILE (Bfile_dir, 'Agent.class');
   ```

   The name of the `libunit` is derived from the name of the class.

Alternatively, you can use the command-line utility `LoadJava`. This uploads Java binaries and resources into a system-generated database table, then uses the `CREATE JAVA` statement to load the Java files into RDBMS libunits. You can upload Java files from file systems, Java IDEs, intranets, or the Internet.

Loading External C Procedures

> **Note:** You can load external C procedures only on platforms that support either DLLs or dynamically loadable shared libraries (such as Solaris `.so` libraries).

When an application calls an external C procedure, Oracle Database or Oracle Listener starts the external procedure agent, `extproc`. Using the network connection established by Oracle Database or Oracle Listener, the application passes this information to `extproc`:

- Name of DLL or shared library
- Name of external procedure

- Any parameters for the external procedure

Then `extproc` loads the DLL or the shared library, runs the external procedure, and passes any values that the external procedure returns back to the application. The application and `extproc` must reside on the same computer.

`extproc` can call procedures in any library that complies with the calling standard used. For more information about the calling standard, see "CALLING STANDARD" on page 14-11.

Note: The default configuration for external procedures no longer requires a network listener to work with Oracle Database and `extproc`. Oracle Database now spawns `extproc` directly, eliminating the risk that Oracle Listener might spawn `extproc` unexpectedly. This default configuration is recommended for maximum security.

You must change this default configuration, so that Oracle Listener spawns `extproc`, if you use any of these:

- A multithreaded `extproc` agent
- Oracle Database in shared mode on Windows
- An `AGENT` clause in the `LIBRARY` specification or an `AGENT IN` clause in the `PROCEDURE` specification that redirects external procedures to a different `extproc` agent

Changing the default configuration requires additional network configuration steps.

To configure your database to use external procedures that are written in C, or that can be called from C applications, you or your database administrator must follow these steps:

1. Define the C Procedures
2. Set Up the Environment
3. Identify the DLL
4. Publish the External Procedures

Define the C Procedures

Define the C procedures using one of these prototypes:

- Kernighan & Ritchie style prototypes; for example:

```
void C_findRoot(x)
 float x;
...
```

- ISO/ANSI prototypes other than numeric data types that are less than full width (such as `float`, `short`, `char`); for example:

```
void C_findRoot(double x)
...
```

- Other data types that do not change size under default argument promotions.

  This example changes size under default argument promotions:

```
void C_findRoot(float x)
...
```

Set Up the Environment

When you use the default configuration for external procedures, Oracle Database spawns `extproc` directly. You need not make configuration changes for `listener.ora` and `tnsnames.ora`. Define the environment variables to be used by external procedures in the file `extproc.ora` (located at `$ORACLE_HOME/hs/admin` on UNIX operating systems and at `ORACLE_HOME\hs\admin` on Windows), using this syntax:

```
SET name=value (environment_variable_name value)
```

Set the `EXTPROC_DLLS` environment variable, which restricts the DLLs that `extproc` can load, to one of these values:

- `NULL`; for example:

  ```
  SET EXTPROC_DLLS=
  ```

  This setting, the default, allows `extproc` to load only the DLLs that are in directory `$ORACLE_HOME/bin` or `$ORACLE_HOME/lib`.

- `ONLY` followed by a colon-separated list of DLLs; for example:

  ```
  SET EXTPROC_DLLS=ONLY:DLL1:DLL2
  ```

  This setting allows `extproc` to load only the DLLs named DLL1 and DLL2. This setting provides maximum security.

- A colon-separated list of DLLs; for example:

  ```
  SET EXTPROC_DLLS=DLL1:DLL2
  ```

  This setting allows `extproc` to load the DLLs named DLL1 and DLL2 and the DLLs that are in directory `$ORACLE_HOME/bin` or `$ORACLE_HOME/lib`.

- `ANY`; for example:

  ```
  SET EXTPROC_DLLS=ANY
  ```

  This setting allows `extproc` to load any DLL.

To change the default configuration for external procedures and have your `extproc` agent spawned by Oracle Listener, configure your database to use external procedures that are written in C, or can be called from C applications, as follows:

1. Set configuration parameters for the agent, named `extproc` by default, in the configuration files `tnsnames.ora` and `listener.ora`. This establishes the connection for the external procedure agent, `extproc`, when the database is started.

2. Start a listener process exclusively for external procedures.

   The Listener sets a few required environment variables (such as `ORACLE_HOME`, `ORACLE_SID`, and `LD_LIBRARY_PATH`) for `extproc`. It can also define specific environment variables in the `ENVS` section of its `listener.ora` entry, and these variables are passed to the agent process. Otherwise, it provides the agent with a "clean" environment. The environment variables set for the agent are independent of those set for the client and server. Therefore, external procedures, which run in the agent process, cannot read environment variables set for the client or server processes.

> **Note:** It is possible for you to set and read environment variables themselves by using the standard C procedures `setenv` and `getenv`, respectively. Environment variables, set this way, are specific to the agent process, which means that they can be read by all functions executed in that process, but not by any other process running on the same host.

3. Determine whether the agent for your external procedure is to run in dedicated mode (the default) or multithreaded mode. In dedicated mode, one "dedicated" agent is launched for each session. In multithreaded mode, a single multithreaded `extproc` agent is launched. The multithreaded `extproc` agent handles calls using different threads for different users. In a configuration where many users can call the external procedures, using a multithreaded `extproc` agent is recommended to conserve system resources.

   If the agent is to run in dedicated mode, additional configuration of the agent process is not necessary.

   If the agent is to run in multithreaded mode, your database administrator must configure the database system to start the agent in multithreaded mode (as a multithreaded `extproc` agent). To do this configuration, use the agent control utility, `agtctl`. For example, start `extproc` using this command:

   ```
   agtctl startup extproc agent_sid
   ```

   where *`agent_sid`* is the system identifier that this `extproc` agent services. An entry for this system identifier is typically added as an entry in the file `tnsnames.ora`. For more information about using `agtctl` for `extproc` administration, see "Administering the Multithreaded extproc Agent" on page A-4.

   > **Note:**
   >
   > - If you use a multithreaded `extproc` agent, the library you call must be thread safe—to avoid errors such as a damaged call stack.
   > - The database server, the agent process, and the listener process that spawns the agent process must all reside on the same host.
   > - By default, the agent process runs on the same database instance as your main application. In situations where reliability is critical, you might want to run the agent process for the external procedure on a separate database instance (still on the same host), so that any problems in the agent do not affect the primary database server. To do so, specify the separate database instance using a database link.

Figure A–1 on page A-3 illustrates the architecture of the multithreaded `extproc` agent.

Identify the DLL

In this context, a DLL is any dynamically loadable operating-system file that stores external procedures.

For security reasons, your DBA controls access to the DLL. Using the `CREATE LIBRARY` statement, the DBA creates a schema object called an alias library, which represents the DLL. Then, if you are an authorized user, the DBA grants you `EXECUTE` privileges on the **alias** library. Alternatively, the DBA might grant you `CREATE ANY`

LIBRARY privileges, in which case you can create your own alias libraries using this syntax:

```
CREATE LIBRARY [schema_name.]library_name
  {IS | AS} 'file_path'
  [AGENT 'agent_link'];
```

> **Note:** The ANY privileges are very powerful and must not be granted lightly. For more information, see:
>
> - *Oracle Database Security Guide* for information about managing system privileges, including ANY
> - *Oracle Database Security Guide* for guidelines for securing user accounts and privileges

It is recommended that you specify the full path to the DLL, rather than just the DLL name. In this example, you create alias library c_utils, which represents DLL utils.so:

```
CREATE LIBRARY C_utils AS '/DLLs/utils.so';
```

To allow flexibility in specifying the DLLs, you can specify the root part of the path as an environment variable using the notation ${*VAR_NAME*}, and set up that variable in the ENVS section of the listener.ora entry.

In this example, the agent specified by the name agent_link is used to run any external procedure in the library C_Utils:

```
create or replace database link agent_link using 'agent_tns_alias';
create or replace library C_utils is
  '${EP_LIB_HOME}/utils.so' agent 'agent_link';
```

The environment variable EP_LIB_HOME is expanded by the agent to the appropriate path for that instance, such as /usr/bin/dll. Variable EP_LIB_HOME must be set in the file listener.ora, for the agent to be able to access it.

For security reasons, extproc, by default, loads only DLLs that are in directory $ORACLE_HOME/bin or $ORACLE_HOME/lib. Also, only local sessions—that is, Oracle Database client processes that run on the same system—are allowed to connect to extproc.

To load DLLs from other directories, set the environment variable EXTPROC_DLLS. The value for this environment variable is a colon-separated list of DLL names qualified with the complete path. For example:

```
EXTPROC_DLLS=/private1/home/johndoe/dll/myDll.so:/private1/home/johndoe/dll/newDll.so
```

While you can set up environment variables for extproc through the ENVS parameter in the file listener.ora, you can also set up environment variables in the extproc initialization file extproc.ora in directory $ORACLE_HOME/hs/admin. When both extproc.ora and ENVS parameter in listener.ora are used, the environment variables defined in extproc.ora take precedence. See the Oracle Net manual for more information about the EXTPROC feature.

Note:

- On a Windows system, specify the path using a drive letter and backslash (\) in the path.
- This technique does not apply to VMS systems, where the ENVS section of listener.ora is not supported.

Publish the External Procedures

You find or write an external C procedure, and add it to the DLL. When the procedure is in the DLL, you publish it using the call specification mechanism described in "Publishing External Procedures" on page 14-9.

Publishing External Procedures

Oracle Database can only use external procedures that are published through a call specification, which maps names, parameter types, and return types for your Java class method or C external procedure to their SQL counterparts. It is written like any other PL/SQL stored procedure except that, in its body, instead of declarations and a BEGIN END block, you code the AS LANGUAGE clause.

The AS LANGUAGE clause specifies:

- Which language the procedure is written in.
- For a Java method:
  - The signature of the Java method.
- For a C procedure:
  - The alias library corresponding to the DLL for a C procedure.
  - The name of the C procedure in a DLL.
  - Various options for specifying how parameters are passed.
  - Which parameter (if any) holds the name of the external procedure agent, extproc, for running the procedure on a different system.

You begin the declaration using the normal CREATE OR REPLACE syntax for a procedure, function, package specification, package body, type specification, or type body.

The call specification follows the name and parameter declarations. Its syntax is:

```
{IS | AS} LANGUAGE {C | JAVA}
```

This is then followed by either:

```
NAME  java_string_literal_name
```

Where *java_string_literal_name* is the signature of your Java method, or by:

```
LIBRARY library_name
[NAME c_string_literal_name]
[WITH CONTEXT]
[PARAMETERS (external_parameter[, external_parameter]...)];
```

Where library_name is the name of your alias library, c_string_literal_name is the name of your external C procedure, and external_parameter stands for:

```
{  CONTEXT
 | SELF [{TDO | property}]
 | {parameter_name | RETURN} [property] [BY REFERENCE] [external_datatype]}
```

`property` stands for:

```
{INDICATOR [{STRUCT | TDO}] | LENGTH | DURATION | MAXLEN | CHARSETID |
CHARSETFORM}
```

> **Note:** Unlike Java, C does not understand SQL types; therefore, the syntax is more intricate

Topics:

- AS LANGUAGE Clause for Java Class Methods
- AS LANGUAGE Clause for External C Procedures

AS LANGUAGE Clause for Java Class Methods

The `AS LANGUAGE` clause is the interface between PL/SQL and a Java class method.

AS LANGUAGE Clause for External C Procedures

These subclauses tell PL/SQL where to locate the external C procedure, how to call it, and what to pass to it:

- LIBRARY
- NAME
- LANGUAGE
- CALLING STANDARD
- WITH CONTEXT
- PARAMETERS
- AGENT IN

Of the preceding subclauses, only `LIBRARY` is required.

LIBRARY

Specifies a local alias library. (You cannot use a database link to specify a remote library.) The library name is a PL/SQL identifier. Therefore, if you enclose the name in double quotation marks, then it becomes case-sensitive. (By default, the name is stored in upper case.) You must have `EXECUTE` privileges on the alias library.

NAME

Specifies the external C procedure to be called. If you enclose the procedure name in double quotation marks, then it becomes case-sensitive. (By default, the name is stored in upper case.) If you omit this subclause, then the procedure name defaults to the upper-case name of the PL/SQL procedure.

> **Note:** The terms `LANGUAGE` and `CALLING STANDARD` apply only to the superseded `AS EXTERNAL` clause.

LANGUAGE

Specifies the third-generation language in which the external procedure was written. If you omit this subclause, then the language name defaults to C.

CALLING STANDARD

Specifies the calling standard under which the external procedure was compiled. The supported calling standard is C. If you omit this subclause, then the calling standard defaults to C.

WITH CONTEXT

Specifies that a context pointer is passed to the external procedure. The context data structure is opaque to the external procedure but is available to service procedures called by the external procedure.

PARAMETERS

Specifies the positions and data types of parameters passed to the external procedure. It can also specify parameter properties, such as current length and maximum length, and the preferred parameter passing method (by value or by reference).

AGENT IN

Specifies which parameter holds the name of the agent process that runs this procedure. This is intended for situations where the external procedure agent, `extproc`, runs using multiple agent processes, to ensure robustness if the agent process of one external procedure fails. You can pass the name of the agent process (corresponding to the name of a database link), and if `tnsnames.ora` and `listener.ora` are set up properly across both instances, the external procedure is called on the other instance. Both instances must be on the same host.

This is similar to the `AGENT` clause of the `CREATE LIBRARY` statement; specifying the value at run time through `AGENT IN` allows greater flexibility.

When the agent name is specified this way, it overrides any agent name declared in the alias library. If no agent name is specified, the default is the `extproc` agent on the same instance as the calling program.

Publishing Java Class Methods

Java classes and their methods are stored in RDBMS libunits in which you can load Java sources, binaries and resources using the `LOADJAVA` utility or the `CREATEJAVA` SQL statements. Libunits can be considered analogous to DLLs written, for example, in C—although they map one-to-one with Java classes, whereas DLLs can contain multiple procedures.

The `NAME`-clause string uniquely identifies the Java method. The PL/SQL function or procedure and Java must have corresponding parameters. If the Java method takes no parameters, then you must code an empty parameter list for it.

When you load Java classes into the RDBMS, they are not published to SQL automatically. This is because the methods of many Java classes are called only from other Java classes, or take parameters for which there is no appropriate SQL type.

Suppose you want to publish this Java method named `J_calcFactorial`, which returns the factorial of its argument:

```
package myRoutines.math;
public class Factorial {
```

```
   public static int J_calcFactorial (int n) {
      if (n == 1) return 1;
      else return n * J_calcFactorial(n - 1);
   }
}
```

This call specification publishes Java method `J_calcFactorial` as PL/SQL stored function `plsToJavaFac_func`, using SQL*Plus:

```
CREATE OR REPLACE FUNCTION Plstojavafac_func (N NUMBER) RETURN NUMBER AS
   LANGUAGE JAVA
   NAME 'myRoutines.math.Factorial.J_calcFactorial(int) return int';
```

Publishing External C Procedures

In this example, you write a PL/SQL standalone function named `plsCallsCdivisor_func` that publishes C function `Cdivisor_func` as an external function:

```
CREATE OR REPLACE FUNCTION Plscallscdivisor_func (
/* Find greatest common divisor of x and y: */
 x     PLS_INTEGER,
 y     PLS_INTEGER)
RETURN PLS_INTEGER
AS LANGUAGE C
   LIBRARY C_utils
   NAME "Cdivisor_func"; /* Quotation marks preserve case. */
```

Locations of Call Specifications

For both Java class methods and external C procedures, call specifications can be specified in any of these locations:

- Standalone PL/SQL procedures
- PL/SQL Package Specifications
- PL/SQL Package Bodies
- ADT Specifications
- ADT Bodies

> **Note:** In Oracle Database version 8.0, `AS EXTERNAL` did not allow call specifications in package or type bodies.

See Also:

- *Oracle Database PL/SQL Language Reference* for more information about calling external procedures from PL/SQL
- *Oracle Database SQL Language Reference* for more information about the SQL `CALL` statement

Examples:

- Example: Locating a Call Specification in a PL/SQL Package
- Example: Locating a Call Specification in a PL/SQL Package Body

- Example: Locating a Call Specification in an ADT Specification
- Example: Locating a Call Specification in an ADT Body
- Example: Java with AUTHID
- Example: C with Optional AUTHID
- Example: Mixing Call Specifications in a Package

> **Note:** In these examples, the `AUTHID` and `SQL_NAME_RESOLVE` clauses might be required to fully stipulate a call specification.

Example: Locating a Call Specification in a PL/SQL Package

```
CREATE OR REPLACE PACKAGE Demo_pack
AUTHID DEFINER
AS
   PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE C
      NAME "C_demoExternal"
      LIBRARY SomeLib
      WITH CONTEXT
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);
END;
```

Example: Locating a Call Specification in a PL/SQL Package Body

```
CREATE OR REPLACE PACKAGE Demo_pack
   AUTHID CURRENT_USER
AS
   PROCEDURE plsToC_demoExternal_proc(x PLS_INTEGER, y VARCHAR2, z DATE);
END;

CREATE OR REPLACE PACKAGE BODY Demo_pack
   SQL_NAME_RESOLVE CURRENT_USER
AS
   PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE JAVA
      NAME 'pkg1.class4.methodProc1(int,java.lang.String,java.sql.Date)';
END;
```

Example: Locating a Call Specification in an ADT Specification

> **Note:** For examples in this topic to work, you must set up this data structure (which requires that you have the privilege `CREATE ANY LIBRARY`):
>
> ```
> CREATE OR REPLACE LIBRARY SOMELIB AS '/tmp/lib.so';
> ```

```
CREATE OR REPLACE TYPE Demo_typ
AUTHID DEFINER
AS OBJECT
   (Attribute1   VARCHAR2(2000), SomeLib varchar2(20),
   MEMBER PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE C
      NAME "C_demoExternal"
      LIBRARY SomeLib
```

```
      WITH CONTEXT
    --  PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE)
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE, SELF)
);
```

Example: Locating a Call Specification in an ADT Body

```
CREATE OR REPLACE TYPE Demo_typ
AUTHID CURRENT_USER
AS OBJECT
   (attribute1 NUMBER,
   MEMBER PROCEDURE plsToJ_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
);

CREATE OR REPLACE TYPE BODY Demo_typ
AS
   MEMBER PROCEDURE plsToJ_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE JAVA
      NAME 'pkg1.class4.J_demoExternal(int,java.lang.String,java.sql.Date)';
END;
```

Example: Java with AUTHID

Here is an example of a publishing a Java class method in a standalone PL/SQL procedure.

```
CREATE OR REPLACE PROCEDURE plsToJ_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z
DATE)
   AUTHID CURRENT_USER
AS LANGUAGE JAVA
   NAME 'pkg1.class4.methodProc1(int,java.lang.String,java.sql.Date)';
```

Example: C with Optional AUTHID

Here is an example of `AS EXTERNAL` publishing a C procedure in a standalone PL/SQL program, in which the `AUTHID` clause is optional. This maintains compatibility with the external procedures of Oracle Database version 8.0.

```
CREATE OR REPLACE PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z
DATE)
AS
   EXTERNAL
   LANGUAGE C
   NAME "C_demoExternal"
   LIBRARY SomeLib
   WITH CONTEXT
   PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);
```

Example: Mixing Call Specifications in a Package

```
CREATE OR REPLACE PACKAGE Demo_pack
AUTHID DEFINER
AS
   PROCEDURE plsToC_InBodyOld_proc (x PLS_INTEGER, y VARCHAR2, z DATE);
   PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE);
   PROCEDURE plsToC_InBody_proc (x PLS_INTEGER, y VARCHAR2, z DATE);
   PROCEDURE plsToJ_InBody_proc (x PLS_INTEGER, y VARCHAR2, z DATE);

   PROCEDURE plsToJ_InSpec_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   IS LANGUAGE JAVA
```

```
      NAME 'pkg1.class4.J_InSpec_meth(int,java.lang.String,java.sql.Date)';

PROCEDURE C_InSpec_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE C
      NAME "C_demoExternal"
      LIBRARY SomeLib
      WITH CONTEXT
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);
END;

CREATE OR REPLACE PACKAGE BODY Demo_pack
AS
PROCEDURE plsToC_InBodyOld_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS EXTERNAL
      LANGUAGE C
      NAME "C_InBodyOld"
      LIBRARY SomeLib
      WITH CONTEXT
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);
PROCEDURE plsToC_demoExternal_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE C
      NAME "C_demoExternal"
      LIBRARY SomeLib
      WITH CONTEXT
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);

PROCEDURE plsToC_InBody_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   AS LANGUAGE C
      NAME "C_InBody"
      LIBRARY SomeLib
      WITH CONTEXT
      PARAMETERS(CONTEXT, x INT, y STRING, z OCIDATE);
PROCEDURE plsToJ_InBody_proc (x PLS_INTEGER, y VARCHAR2, z DATE)
   IS LANGUAGE JAVA
      NAME 'pkg1.class4.J_InBody_meth(int,java.lang.String,java.sql.Date)';
END;
```

Passing Parameters to External C Procedures with Call Specifications

Call specifications allows a mapping between PL/SQL and C data types. See Specifying Data Types for data type mappings.

Passing parameters to an external C procedure is complicated by several circumstances:

- The available set of PL/SQL data types does not correspond one-to-one with the set of C data types.
- Unlike C, PL/SQL includes the RDBMS concept of nullity. Therefore, PL/SQL parameters can be `NULL`, whereas C parameters cannot.
- The external procedure might need the current length or maximum length of `CHAR`, `LONG RAW`, `RAW`, and `VARCHAR2` parameters.
- The external procedure might need character set information about `CHAR`, `VARCHAR2`, and `CLOB` parameters.
- PL/SQL might need the current length, maximum length, or null status of values returned by the external procedure.

Note: The maximum number of parameters that you can pass to a C external procedure is 128. However, if you pass float or double parameters by value, then the maximum is less than 128. How much less depends on the number of such parameters and your operating system. To get a rough estimate, count each float or double passed by value as two parameters.

Topics:

- Specifying Data Types
- External Data Type Mappings
- Passing Parameters BY VALUE or BY REFERENCE
- Declaring Formal Parameters
- Overriding Default Data Type Mapping
- Specifying Properties

Specifying Data Types

Do not pass parameters to an external procedure directly. Instead, pass them to the PL/SQL procedure that published the external procedure, specifying PL/SQL data types for the parameters. PL/SQL data types map to default external data types, as shown in Table 14–1.

Note: The PL/SQL data types `BINARY_INTEGER` and `PLS_INTEGER` are identical. For simplicity, this document uses "`PLS_INTEGER`" to mean both `BINARY_INTEGER` and `PLS_INTEGER`.

Table 14–1 Parameter Data Type Mappings

PL/SQL Data Type	Supported External Types	Default External Type
BINARY_INTEGER BOOLEAN PLS_INTEGER	[UNSIGNED] CHAR [UNSIGNED] SHORT [UNSIGNED] INT [UNSIGNED] LONG SB1, SB2, SB4 UB1, UB2, UB4 SIZE_T	INT
NATURAL[1] NATURALN[1] POSITIVE[1] POSITIVEN[1] SIGNTYPE[1]	[UNSIGNED] CHAR [UNSIGNED] SHORT [UNSIGNED] INT [UNSIGNED] LONG SB1, SB2, SB4 UB1, UB2, UB4 SIZE_T	UNSIGNED INT
FLOAT REAL	FLOAT	FLOAT
DOUBLE PRECISION	DOUBLE	DOUBLE

Table 14–1 (Cont.) Parameter Data Type Mappings

PL/SQL Data Type	Supported External Types	Default External Type
CHAR CHARACTER LONG NCHAR NVARCHAR2 ROWID VARCHAR VARCHAR2	STRING OCISTRING	STRING
LONG RAW RAW	RAW OCIRAW	RAW
BFILE BLOB CLOB NCLOB	OCILOBLOCATOR	OCILOBLOCATOR
NUMBER DEC[1] DECIMAL[1] INT[1] INTEGER[1] NUMERIC[1] SMALLINT[1]	OCINUMBER	OCINUMBER
DATE	OCIDATE	OCIDATE
TIMESTAMP TIMESTAMP WITH TIME ZONE TIMESTAMP WITH LOCAL TIME ZONE	OCIDateTime	OCIDateTime
INTERVAL DAY TO SECOND INTERVAL YEAR TO MONTH	OCIInterval	OCIInterval
composite object types: ADTs	dvoid	dvoid
composite object types: collections (varrays, nested tables)	OCICOLL	OCICOLL

[1] This PL/SQL type compiles only if you use AS EXTERNAL in your call spec.

External Data Type Mappings

Each external data type maps to a C data type, and the data type conversions are performed implicitly. To avoid errors when declaring C prototype parameters, see Table 14–2, which shows the C data type to specify for a given external data type and PL/SQL parameter mode. For example, if the external data type of an OUT parameter is STRING, then specify the data type char * in your C prototype.

Table 14–2 External Data Type Mappings

External Data Type Corresponding to PL/SL Type	If Mode is IN or RETURN, Specify in C Prototype...	If Mode is IN by Reference or RETURN by Reference, Specify in C Prototype...	If Mode is IN OUT or OUT, Specify in C Prototype...
CHAR	char	char *	char *
UNSIGNED CHAR	unsigned char	unsigned char *	unsigned char *
SHORT	short	short *	short *
UNSIGNED SHORT	unsigned short	unsigned short *	unsigned short *

Table 14–2 (Cont.) External Data Type Mappings

External Data Type Corresponding to PL/SL Type	If Mode is IN or RETURN, Specify in C Prototype...	If Mode is IN by Reference or RETURN by Reference, Specify in C Prototype...	If Mode is IN OUT or OUT, Specify in C Prototype...
INT	int	int *	int *
UNSIGNED INT	unsigned int	unsigned int *	unsigned int *
LONG	long	long *	long *
UNSIGNED LONG	unsigned long	unsigned long *	unsigned long *
CHAR	char	char *	char *
UNSIGNED CHAR	unsigned char	unsigned char *	unsigned char *
SHORT	short	short *	short *
UNSIGNED SHORT	unsigned short	unsigned short *	unsigned short *
INT	int	int *	int *
UNSIGNED INT	unsigned int	unsigned int *	unsigned int *
LONG	long	long *	long *
UNSIGNED LONG	unsigned long	unsigned long *	unsigned long *
SIZE_T	size_t	size_t *	size_t *
SB1	sb1	sb1 *	sb1 *
UB1	ub1	ub1 *	ub1 *
SB2	sb2	sb2 *	sb2 *
UB2	ub2	ub2 *	ub2 *
SB4	sb4	sb4 *	sb4 *
UB4	ub4	ub4 *	ub4 *
FLOAT	float	float *	float *
DOUBLE	double	double *	double *
STRING	char *	char *	char *
RAW	unsigned char *	unsigned char *	unsigned char *
OCILOBLOCATOR	OCILobLocator *	OCILobLocator **	OCILobLocator **
OCINUMBER	OCINumber *	OCINumber *	OCINumber *
OCISTRING	OCIString *	OCIString *	OCIString *
OCIRAW	OCIRaw *	OCIRaw *	OCIRaw *
OCIDATE	OCIDate *	OCIDate *	OCIDate *
OCICOLL	OCIColl * or OCIArray * or OCITable *	OCIColl ** or OCIArray ** or OCITable **	OCIColl ** or OCIArray ** or OCITable **
OCITYPE	OCIType *	OCIType *	OCIType *
TDO	OCIType *	OCIType *	OCIType *
ADT (final types)	dvoid*	dvoid*	dvoid*
ADT (nonfinal types)	dvoid*	dvoid*	dvoid**

Composite data types are not self describing. Their description is stored in a **Type Descriptor Object** (TDO). Objects and indicator structs for objects have no predefined OCI data type, but must use the data types generated by Oracle Database's **Object Type Translator** (OTT). The optional TDO argument for `INDICATOR`, and for composite objects, in general, has the C data type, OCIType *.

`OCICOLL` for `REF` and collection arguments is optional and exists only for completeness. You cannot map a `REF` or collection type onto any other data type, or any other data type onto a `REF` or collection type.

Passing Parameters BY VALUE or BY REFERENCE

If you specify `BY VALUE`, then scalar `IN` and `RETURN` arguments are passed by value (which is also the default). Alternatively, you might have them passed by reference by specifying `BY REFERENCE`.

By default, or if you specify `BY REFERENCE`, then scalar `IN OUT`, and `OUT` arguments are passed by reference. Specifying `BY VALUE` for `IN OUT`, and `OUT` arguments is not supported for C. The usefulness of the `BY REFERENCE/VALUE` clause is restricted to external data types that are, by default, passed by value. This is true for `IN`, and `RETURN` arguments of these external types:

```
[UNSIGNED] CHAR
[UNSIGNED] SHORT
[UNSIGNED] INT
[UNSIGNED] LONG
SIZE_T
SB1
SB2
SB4
UB1
UB2
UB4
FLOAT
DOUBLE
```

All `IN` and `RETURN` arguments of external types not on this list, all `IN OUT` arguments, and all `OUT` arguments are passed by reference.

Declaring Formal Parameters

Generally, the PL/SQL procedure that publishes an external procedure declares a list of formal parameters, as this example shows:

> **Note:** You might need to set up this data structure for examples in this topic to work:
>
> ```
> CREATE LIBRARY MathLib AS '/tmp/math.so';
> ```

```
CREATE OR REPLACE FUNCTION Interp_func (
/* Find the value of y at x degrees using Lagrange interpolation: */
   x    IN FLOAT,
   y    IN FLOAT)
RETURN FLOAT AS
   LANGUAGE C
   NAME "Interp_func"
   LIBRARY MathLib;
```

Each formal parameter declaration specifies a name, parameter mode, and PL/SQL data type (which maps to the default external data type). That might be all the information the external procedure needs. If not, then you can provide more information using the PARAMETERS clause, which lets you specify:

- Nondefault external data types
- The current or maximum length of a parameter
- NULL/NOT NULL indicators for parameters
- Character set IDs and forms
- The position of parameters in the list
- How IN parameters are passed (by value or by reference)

If you decide to use the PARAMETERS clause, keep in mind:

- For every formal parameter, there must be a corresponding parameter in the PARAMETERS clause.
- If you include the WITH CONTEXT clause, then you must specify the parameter CONTEXT, which shows the position of the context pointer in the parameter list.
- If the external procedure is a function, then you might specify the RETURN parameter, but it must be in the last position. If RETURN is not specified, the default external type is used.

Overriding Default Data Type Mapping

In some cases, you can use the PARAMETERS clause to override the default data type mappings. For example, you can remap the PL/SQL data type BOOLEAN from external data type INT to external data type CHAR.

Specifying Properties

You can also use the PARAMETERS clause to pass more information about PL/SQL formal parameters and function results to an external procedure. Do this by specifying one or more of these properties:

```
INDICATOR [{STRUCT | TDO}]
LENGTH
DURATION
MAXLEN
CHARSETID
CHARSETFORM
SELF
```

Table 14–3 shows the allowed and the default external data types, PL/SQL data types, and PL/SQL parameter modes allowed for a given property. MAXLEN (used to specify data returned from C back to PL/SQL) cannot be applied to an IN parameter.

Table 14–3 Properties and Data Types

Property	Allowed External Types (C)	Default External Type (C)	Allowed PL/SQL Types	Allowed PL/SQL Modes	Default PL/SQL Passing Method
INDICATOR	SHORT	SHORT	all scalars	IN IN OUT OUT RETURN	BY VALUE BY REFERENCE BY REFERENCE BY REFERENCE
LENGTH	[UNSIGNED] SHORT [UNSIGNED] INT [UNSIGNED] LONG	INT	CHAR LONG RAW RAW VARCHAR2	IN IN OUT OUT RETURN	BY VALUE BY REFERENCE BY REFERENCE BY REFERENCE
MAXLEN	[UNSIGNED] SHORT [UNSIGNED] INT [UNSIGNED] LONG	INT	CHAR LONG RAW RAW VARCHAR2	IN OUT OUT RETURN	BY REFERENCE BY REFERENCE BY REFERENCE
CHARSETID CHARSETFORM	UNSIGNED SHORT UNSIGNED INT UNSIGNED LONG	UNSIGNED INT	CHAR CLOB VARCHAR2	IN IN OUT OUT RETURN	BY VALUE BY REFERENCE BY REFERENCE BY REFERENCE

In this example, the PARAMETERS clause specifies properties for the PL/SQL formal parameters and function result:

```
CREATE OR REPLACE FUNCTION plsToCparse_func  (
   x   IN PLS_INTEGER,
   Y   IN OUT CHAR)
RETURN CHAR AS LANGUAGE C
   LIBRARY c_utils
   NAME "C_parse"
   PARAMETERS (
      x,            -- stores value of x
      x INDICATOR,  -- stores null status of x
      y,            -- stores value of y
      y LENGTH,     -- stores current length of y
      y MAXLEN,     -- stores maximum length of y
      RETURN INDICATOR,
      RETURN);
```

With this PARAMETERS clause, the C prototype becomes:

```
char  *C_parse( int x, short x_ind, char *y, int *y_len, int *y_maxlen,
  short *retind );
```

The additional parameters in the C prototype correspond to the INDICATOR (for x), LENGTH (of y), and MAXLEN (of y), and the INDICATOR for the function result in the PARAMETERS clause. The parameter RETURN corresponds to the C function identifier, which stores the result value.

Topics:

- INDICATOR
- LENGTH and MAXLEN
- CHARSETID and CHARSETFORM
- Repositioning Parameters
- SELF

- BY REFERENCE
- WITH CONTEXT
- Interlanguage Parameter Mode Mappings

INDICATOR

An INDICATOR is a parameter whose value indicates whether another parameter is NULL. PL/SQL does not need indicators, because the RDBMS concept of nullity is built into the language. However, an external procedure might need to know if a parameter or function result is NULL. Also, an external procedure might need to signal the server that a returned value is actually a NULL, and must be treated accordingly.

In such cases, you can use the property INDICATOR to associate an indicator with a formal parameter. If the PL/SQL procedure is a function, then you can also associate an indicator with the function result, as shown earlier.

To check the value of an indicator, you can use the constants OCI_IND_NULL and OCI_IND_NOTNULL. If the indicator equals OCI_IND_NULL, then the associated parameter or function result is NULL. If the indicator equals OCI_IND_NOTNULL, then the parameter or function result is not NULL.

For IN parameters, which are inherently read-only, INDICATOR is passed by value (unless you specify BY REFERENCE) and is read-only (even if you specify BY REFERENCE). For OUT, IN OUT, and RETURN parameters, INDICATOR is passed by reference by default.

The INDICATOR can also have a STRUCT or TDO option. Because specifying INDICATOR as a property of an object is not supported, and because arguments of objects have complete indicator structs instead of INDICATOR scalars, you must specify this by using the STRUCT option. You must use the type descriptor object (TDO) option for composite objects and collections,

LENGTH and MAXLEN

In PL/SQL, there is no standard way to indicate the length of a RAW or string parameter. However, you often want to pass the length of such a parameter to and from an external procedure. Using the properties LENGTH and MAXLEN, you can specify parameters that store the current length and maximum length of a formal parameter.

> **Note:** With a parameter of type RAW or LONG RAW, you must use the property LENGTH. Also, if that parameter is IN OUT and NULL or OUT and NULL, then you must set the length of the corresponding C parameter to zero.

For IN parameters, LENGTH is passed by value (unless you specify BY REFERENCE) and is read-only. For OUT, IN OUT, and RETURN parameters, LENGTH is passed by reference.

As mentioned earlier, MAXLEN does not apply to IN parameters. For OUT, IN OUT, and RETURN parameters, MAXLEN is passed by reference and is read-only.

CHARSETID and CHARSETFORM

Oracle Database provides globalization support, which lets you process single-byte and multibyte character data and convert between character sets. It also lets your applications run in different language environments.

By default, if the server and agent use the exact same $ORACLE_HOME value, the agent uses the same globalization support settings as the server (including any settings that were specified with ALTER SESSION statements).

If the agent is running in a separate $ORACLE_HOME (even if the same location is specified by two different aliases or symbolic links), the agent uses the same globalization support settings as the server except for the character set; the default character set for the agent is defined by the NLS_LANG and NLS_NCHAR environment settings for the agent.

The properties CHARSETID and CHARSETFORM identify the nondefault character set from which the character data being passed was formed. With CHAR, CLOB, and VARCHAR2 parameters, you can use CHARSETID and CHARSETFORM to pass the character set ID and form to the external procedure.

For IN parameters, CHARSETID and CHARSETFORM are passed by value (unless you specify BY REFERENCE) and are read-only (even if you specify BY REFERENCE). For OUT, IN OUT, and RETURN parameters, CHARSETID and CHARSETFORM are passed by reference and are read-only.

The OCI attribute names for these properties are OCI_ATTR_CHARSET_ID and OCI_ATTR_CHARSET_FORM.

See Also:

- *Oracle Call Interface Programmer's Guide* for more information about OCI_ATTR_CHARSET_ID and OCI_ATTR_CHARSET_FORM
- *Oracle Database Globalization Support Guide* for more information about using national language data with the OCI

Repositioning Parameters

Remember, each formal parameter of the external procedure must have a corresponding parameter in the PARAMETERS clause. Their positions can differ, because PL/SQL associates them by name, not by position. However, the PARAMETERS clause and the C prototype for the external procedure must have the same number of parameters, and they must be in the same order.

SELF

SELF is the always-present argument of an object type's member procedure, namely the object instance itself. In most cases, this argument is implicit and is not listed in the argument list of the PL/SQL procedure. However, SELF must be explicitly specified as an argument of the PARAMETERS clause.

For example, assume that a user wants to create a Person object, consisting of a person's name and date of birth, and then create a table of this object type. The user eventually wants to determine the age of each Person object in this table.

In SQL*Plus, the Person object type can be created by:

```
CREATE OR REPLACE TYPE Person1_typ AS OBJECT (
  Name_     VARCHAR2(30),
  B_date    DATE,
  MEMBER FUNCTION calcAge_func RETURN NUMBER
);
/
```

Typically, the member function is implemented in PL/SQL, but in this example it is an external procedure. The body of the member function is declared as follows:

```
CREATE OR REPLACE TYPE BODY Person1_typ AS
  MEMBER FUNCTION calcAge_func RETURN NUMBER
  AS LANGUAGE C
  NAME "age"
  LIBRARY agelib
  WITH CONTEXT
  PARAMETERS (
    CONTEXT,
    SELF,
    SELF INDICATOR STRUCT,
    SELF TDO,
    RETURN INDICATOR
  );
END;
/
```

The `calcAge_func` member function does not take any arguments, but only returns a number. A member function is always called on an instance of the associated object type. The object instance itself always is an implicit argument of the member function. To refer to the implicit argument, the `SELF` keyword is used. This is incorporated into the external procedure syntax by supporting references to `SELF` in the parameters clause.

The matching table is created and populated.

```
CREATE TABLE Person_tab OF Person1_typ;

INSERT INTO Person_tab
VALUES ('BOB', TO_DATE('14-MAY-85'));

INSERT INTO Person_tab
VALUES ('JOHN', TO_DATE('22-DEC-71'));
```

Finally, retrieve the information of interest from the table.

```
SELECT p.name, p.b_date, p.calcAge_func() FROM Person_tab p;

NAME                           B_DATE    P.CALCAGE_
------------------------------ --------- ----------
BOB                            14-MAY-85          0
JOHN                           22-DEC-71          0
```

This is sample C code that implements the `external` member function and the Object-Type-Translator (OTT)-generated `struct` definitions:

```
#include <oci.h>

struct PERSON
{
    OCIString   *NAME;
    OCIDate      B_DATE;
};
typedef struct PERSON PERSON;

struct PERSON_ind
{
    OCIInd    _atomic;
    OCIInd    NAME;
```

```
    OCIInd    B_DATE;
};
typedef struct PERSON_ind PERSON_ind;

OCINumber *age (ctx, person_obj, person_obj_ind, tdo, ret_ind)
OCIExtProcContext *ctx;
PERSON         *person_obj;
PERSON_ind     *person_obj_ind;
OCIType        *tdo;
OCIInd         *ret_ind;
{
    sword      err;
    text       errbuf[512];
    OCIEnv    *envh;
    OCISvcCtx *svch;
    OCIError  *errh;
    OCINumber *age;
    int        inum = 0;
    sword      status;

    /* get OCI Environment */
    err = OCIExtProcGetEnv( ctx, &envh, &svch, &errh );

    /* initialize return age to 0 */
    age = (OCINumber *)OCIExtProcAllocCallMemory(ctx, sizeof(OCINumber));
    status = OCINumberFromInt(errh, &inum, sizeof(inum), OCI_NUMBER_SIGNED,
                              age);
    if (status != OCI_SUCCESS)
    {
      OCIExtProcRaiseExcp(ctx, (int)1476);
      return (age);
    }

    /* return NULL if the person object is null or the birthdate is null */
    if ( person_obj_ind->_atomic == OCI_IND_NULL ||
         person_obj_ind->B_DATE  == OCI_IND_NULL )
    {
        *ret_ind = OCI_IND_NULL;
        return (age);
    }

    /* The actual implementation to calculate the age is left to the reader,
       but an easy way of doing this is a callback of the form:
            select trunc(months_between(sysdate, person_obj->b_date) / 12)
            from DUAL;
    */
    *ret_ind = OCI_IND_NOTNULL;
    return (age);
}
```

BY REFERENCE

In C, you can pass `IN` scalar parameters by value (the value of the parameter is passed) or by reference (a pointer to the value is passed). When an external procedure expects a pointer to a scalar, specify `BY REFERENCE` phrase to pass the parameter by reference:

```
CREATE OR REPLACE PROCEDURE findRoot_proc (
   x IN DOUBLE PRECISION)
AS LANGUAGE C
   LIBRARY c_utils
   NAME "C_findRoot"
```

```
   PARAMETERS (
      x BY REFERENCE);
```

In this case, the C prototype is:

```
void C_findRoot(double *x);
```

The default (used when there is no PARAMETERS clause) is:

```
void C_findRoot(double x);
```

WITH CONTEXT

By including the WITH CONTEXT clause, you can give an external procedure access to information about parameters, exceptions, memory allocation, and the user environment. The WITH CONTEXT clause specifies that a context pointer is passed to the external procedure. For example, if you write this PL/SQL function:

```
CREATE OR REPLACE FUNCTION getNum_func (
   x IN REAL)
RETURN PLS_INTEGER AS LANGUAGE C
   LIBRARY c_utils
   NAME "C_getNum"
   WITH CONTEXT
   PARAMETERS (
      CONTEXT,
      x BY REFERENCE,
      RETURN INDICATOR);
```

The C prototype is:

```
int C_getNum(
   OCIExtProcContext *with_context,
   float *x,
   short *retind);
```

The context data structure is opaque to the external procedure; but, is available to service procedures called by the external procedure.

If you also include the PARAMETERS clause, then you must specify the parameter CONTEXT, which shows the position of the context pointer in the parameter list. If you omit the PARAMETERS clause, then the context pointer is the first parameter passed to the external procedure.

Interlanguage Parameter Mode Mappings

PL/SQL supports the IN, IN OUT, and OUT parameter modes, and the RETURN clause for procedures returning values.

Running External Procedures with CALL Statements

Now that you have published your Java class method or external C procedure, you are ready to call it.

Do not call an external procedure directly. Instead, use the CALL statement to call the PL/SQL procedure that published the external procedure. See "CALL Statement Syntax" on page 14-28.

Such calls, which you code in the same manner as a call to a regular PL/SQL procedure, can appear in:

- Anonymous blocks

- Standalone and package procedures
- Methods of an object type
- Database triggers
- SQL statements (calls to package functions only).

Any PL/SQL block or procedure running on the server side, or on the client side, (for example, in a tool such as Oracle Forms) can call an external procedure. On the server side, the external procedure runs in a separate process address space, which safeguards your database. Figure 14–1 shows how Oracle Database and external procedures interact.

Figure 14–1 Oracle Database and External Procedures

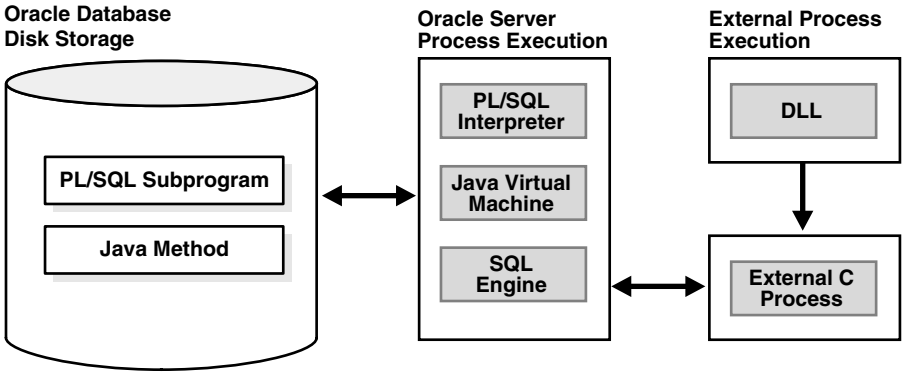

Topics:

- Preconditions for External Procedures
- CALL Statement Syntax
- Calling Java Class Methods
- Calling External C Procedures

Preconditions for External Procedures

Before calling external procedures, consider the privileges, permissions, and synonyms that exist in the execution environment.

Topics:

- Privileges of External Procedures
- Managing Permissions
- Creating Synonyms for External Procedures

Privileges of External Procedures

When external procedures are called through `CALL` specifications, they run with definer's privileges, rather than invoker privileges.

A program running with invoker privileges is not bound to a particular schema. It runs at the calling site and accesses database items (such as tables and views) with the caller's visibility and permissions. However, a program running with definer's privileges is bound to the schema in which it is defined. It runs at the defining site, in

the definer's schema, and accesses database items with the definer's visibility and permissions.

Managing Permissions

To call an external procedure, a user must have the EXECUTE privilege on its call specification. To grant this privilege, use the SQL statement GRANT (described in *Oracle Database SQL Language Reference*). For example, this statement allows the user johndoe to call the external procedure whose call specification is plsToJ_demoExternal_proc:

```
GRANT EXECUTE ON plsToJ_demoExternal_proc TO johndoe;
```

Grant the EXECUTE privilege on a call specification only to users who need to call the procedure.

Creating Synonyms for External Procedures

For convenience, you or your DBA can create synonyms for external procedures using the CREATE PUBLIC SYNONYM statement. In this example, your DBA creates a public synonym, which is accessible to all users. If PUBLIC is not specified, then the synonym is private and accessible only within its schema.

```
CREATE PUBLIC SYNONYM Rfac FOR johndoe.RecursiveFactorial;
```

CALL Statement Syntax

Call the external procedure through the SQL CALL statement. You can run the CALL statement interactively from SQL*Plus. The syntax is:

```
CALL [schema.][{object_type_name | package_name}]procedure_name[@dblink_name]
   [(parameter_list)] [INTO :host_variable][INDICATOR][:indicator_variable];
```

This is equivalent to running a procedure myproc using a SQL statement of the form "SELECT myproc(...) FROM DUAL," except that the overhead associated with performing the SELECT is not incurred.

For example, here is an anonymous PL/SQL block that uses dynamic SQL to call plsToC_demoExternal_proc, which you published. PL/SQL passes three parameters to the external C procedure C_demoExternal_proc.

```
DECLARE
   xx NUMBER(4);
   yy VARCHAR2(10);
   zz DATE;
 BEGIN
 EXECUTE IMMEDIATE
 'CALL plsToC_demoExternal_proc(:xxx, :yyy, :zzz)' USING xx,yy,zz;
 END;
```

The semantics of the CALL statement is identical to the that of an equivalent BEGIN END block.

> **Note:** CALL is the only SQL statement that cannot be put, by itself, in a PL/SQL BEGIN END block. It can be part of an EXECUTE IMMEDIATE statement within a BEGIN END block.

Calling Java Class Methods

To call the `J_calcFactorial` class method published earlier:

1. Declare and initialize two SQL*Plus host variables:

```
VARIABLE x NUMBER
VARIABLE y NUMBER
EXECUTE :x := 5;
```

2. Call `J_calcFactorial`:

```
CALL J_calcFactorial(:x) INTO :y;
PRINT y
```

Result:

```
Y
------
   120
```

Calling External C Procedures

To call an external C procedure, PL/SQL must find the path of the appropriate DLL. The PL/SQL engine retrieves the path from the data dictionary, based on the library alias from the `AS LANGUAGE` clause of the procedure declaration.

Next, PL/SQL alerts a Listener process which, in turn, spawns a session-specific agent. By default, this agent is named `extproc`, although you can specify other names in the `listener.ora` file. The Listener hands over the connection to the agent, and PL/SQL passes to the agent the name of the DLL, the name of the external procedure, and any parameters.

Then, the agent loads the DLL and runs the external procedure. Also, the agent handles service calls (such as raising an exception) and callbacks to Oracle Database. Finally, the agent passes to PL/SQL any values returned by the external procedure.

> **Note:** Although some DLL caching takes place, your DLL might not remain in the cache; therefore, do not store global variables in your DLL.

After the external procedure completes, the agent remains active throughout your Oracle Database session; when you log off, the agent is stopped. Consequently, you incur the cost of launching the agent only once, no matter how many calls you make. Still, call an external procedure only when the computational benefits outweigh the cost.

Here, you call PL/SQL function `plsCallsCdivisor_func`, which you published previously, from an anonymous block. PL/SQL passes the two integer parameters to external function `Cdivisor_func`, which returns their greatest common divisor.

```
DECLARE
   g    PLS_INTEGER;
   a    PLS_INTEGER;
   b    PLS_INTEGER;
CALL plsCallsCdivisor_func(a, b);
IF g IN (2,4,8) THEN ...
```

Handling Errors and Exceptions in Multilanguage Programs

The PL/SQL compiler raises compile-time exceptions if an `AS EXTERNAL` call specification is found in a `TYPE` or `PACKAGE` specification.

C programs can raise exceptions through the `OCIExtproc` functions.

Using Service Routines with External C Procedures

When called from an external procedure, a **service routine** can raise exceptions, allocate memory, and call OCI handles for callbacks to the server. To use a service routine, you must specify the `WITH CONTEXT` clause, which lets you pass a context structure to the external procedure. The context structure is declared in header file `ociextp.h` as follows:

```
typedef struct OCIExtProcContext OCIExtProcContext;
```

> **Note:** `ociextp.h` is located in `$ORACLE_HOME/plsql/public` on Linux and UNIX.

Service procedures:

- OCIExtProcAllocCallMemory
- OCIExtProcRaiseExcp
- OCIExtProcRaiseExcpWithMsg

OCIExtProcAllocCallMemory

This service routine allocates *n* bytes of memory for the duration of the external procedure call. Any memory allocated by the function is freed automatically as soon as control returns to PL/SQL.

> **Note:** Do not have the external procedure call the C function `free` to free memory allocated by this service routine, as this is handled automatically.

The C prototype for this function is as follows:

```
dvoid *OCIExtProcAllocCallMemory(
   OCIExtProcContext *with_context,
   size_t amount);
```

The parameters `with_context` and `amount` are the context pointer and number of bytes to allocate, respectively. The function returns an untyped pointer to the allocated memory. A return value of zero indicates failure.

In SQL*Plus, suppose you publish external function `plsToC_concat_func`, as follows:

```
CREATE OR REPLACE FUNCTION plsToC_concat_func (
   str1 IN VARCHAR2,
   str2 IN VARCHAR2)
RETURN VARCHAR2 AS LANGUAGE C
NAME "concat"
LIBRARY stringlib
WITH CONTEXT
```

```
PARAMETERS (
CONTEXT,
str1   STRING,
str1   INDICATOR short,
str2   STRING,
str2   INDICATOR short,
RETURN INDICATOR short,
RETURN LENGTH short,
RETURN STRING);
```

When called, `C_concat` concatenates two strings, then returns the result:

```
select plsToC_concat_func('hello ', 'world') from DUAL;
PLSTOC_CONCAT_FUNC('HELLO','WORLD')
--------------------------------------------------------------------------------
hello world
```

If either string is `NULL`, the result is also `NULL`. As this example shows, `C_concat` uses `OCIExtProcAllocCallMemory` to allocate memory for the result string:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <oci.h>
#include <ociextp.h>

char *concat(ctx, str1, str1_i, str2, str2_i, ret_i, ret_l)
OCIExtProcContext *ctx;
char   *str1;
short  str1_i;
char   *str2;
short  str2_i;
short  *ret_i;
short  *ret_l;
{
  char *tmp;
  short len;
  /* Check for null inputs. */
  if ((str1_i == OCI_IND_NULL) || (str2_i == OCI_IND_NULL))
  {
      *ret_i = (short)OCI_IND_NULL;
      /* PL/SQL has no notion of a NULL ptr, so return a zero-byte string. */
      tmp = OCIExtProcAllocCallMemory(ctx, 1);
      tmp[0] = '\0';
      return(tmp);
  }
  /* Allocate memory for result string, including NULL terminator. */
  len = strlen(str1) + strlen(str2);
  tmp = OCIExtProcAllocCallMemory(ctx, len + 1);

  strcpy(tmp, str1);
  strcat(tmp, str2);

  /* Set NULL indicator and length. */
  *ret_i = (short)OCI_IND_NOTNULL;
  *ret_l = len;
  /* Return pointer, which PL/SQL frees later. */
  return(tmp);
}

#ifdef LATER
```

```
static void checkerr (/*_ OCIError *errhp, sword status _*/);

void checkerr(errhp, status)
OCIError *errhp;
sword status;
{
  text errbuf[512];
  sb4 errcode = 0;

  switch (status)
  {
  case OCI_SUCCESS:
    break;
  case OCI_SUCCESS_WITH_INFO:
    (void) printf("Error - OCI_SUCCESS_WITH_INFO\n");
    break;
  case OCI_NEED_DATA:
    (void) printf("Error - OCI_NEED_DATA\n");
    break;
  case OCI_NO_DATA:
    (void) printf("Error - OCI_NODATA\n");
    break;
  case OCI_ERROR:
    (void) OCIErrorGet((dvoid *)errhp, (ub4) 1, (text *) NULL, &errcode,
                       errbuf, (ub4) sizeof(errbuf), OCI_HTYPE_ERROR);
    (void) printf("Error - %.*s\n", 512, errbuf);
    break;
  case OCI_INVALID_HANDLE:
    (void) printf("Error - OCI_INVALID_HANDLE\n");
    break;
  case OCI_STILL_EXECUTING:
    (void) printf("Error - OCI_STILL_EXECUTE\n");
    break;
  case OCI_CONTINUE:
    (void) printf("Error - OCI_CONTINUE\n");
    break;
  default:
    break;
  }
}

char *concat(ctx, str1, str1_i, str2, str2_i, ret_i, ret_l)
OCIExtProcContext *ctx;
char   *str1;
short  str1_i;
char   *str2;
short  str2_i;
short  *ret_i;
short  *ret_l;
{
  char *tmp;
  short len;
  /* Check for null inputs. */
  if ((str1_i == OCI_IND_NULL) || (str2_i == OCI_IND_NULL))
  {
      *ret_i = (short)OCI_IND_NULL;
      /* PL/SQL has no notion of a NULL ptr, so return a zero-byte string. */
      tmp = OCIExtProcAllocCallMemory(ctx, 1);
      tmp[0] = '\0';
      return(tmp);
```

```
  }
  /* Allocate memory for result string, including NULL terminator. */
  len = strlen(str1) + strlen(str2);
  tmp = OCIExtProcAllocCallMemory(ctx, len + 1);

  strcpy(tmp, str1);
  strcat(tmp, str2);

  /* Set NULL indicator and length. */
  *ret_i = (short)OCI_IND_NOTNULL;
  *ret_l = len;
  /* Return pointer, which PL/SQL frees later. */
  return(tmp);
}

/*======================================================================*/
int main(char *argv, int argc)
{
  OCIExtProcContext *ctx;
  char           *str1;
  short          str1_i;
  char           *str2;
  short          str2_i;
  short          *ret_i;
  short          *ret_l;
  /* OCI Handles */
  OCIEnv        *envhp;
  OCIServer     *srvhp;
  OCISvcCtx     *svchp;
  OCIError      *errhp;
  OCISession    *authp;
  OCIStmt       *stmthp;
  OCILobLocator *clob, *blob;
  OCILobLocator *Lob_loc;

  /* Initialize and Logon */
  (void) OCIInitialize((ub4) OCI_DEFAULT, (dvoid *)0,
                       (dvoid * (*)(dvoid *, size_t)) 0,
                       (dvoid * (*)(dvoid *, dvoid *, size_t))0,
                       (void (*)(dvoid *, dvoid *)) 0 );

  (void) OCIEnvInit( (OCIEnv **) &envhp,
                    OCI_DEFAULT, (size_t) 0,
                    (dvoid **) 0 );

  (void) OCIHandleAlloc( (dvoid *) envhp, (dvoid **) &errhp, OCI_HTYPE_ERROR,
                   (size_t) 0, (dvoid **) 0);

  /* Server contexts */
  (void) OCIHandleAlloc( (dvoid *) envhp, (dvoid **) &srvhp, OCI_HTYPE_SERVER,
                   (size_t) 0, (dvoid **) 0);

  /* Service context */
  (void) OCIHandleAlloc( (dvoid *) envhp, (dvoid **) &svchp, OCI_HTYPE_SVCCTX,
                   (size_t) 0, (dvoid **) 0);

  /* Attach to Oracle Database */
  (void) OCIServerAttach( srvhp, errhp, (text *)"", strlen(""), 0);

  /* Set attribute server context in the service context */
```

```
  (void) OCIAttrSet ((dvoid *) svchp, OCI_HTYPE_SVCCTX,
                     (dvoid *)srvhp, (ub4) 0,
                    OCI_ATTR_SERVER, (OCIError *) errhp);

  (void) OCIHandleAlloc((dvoid *) envhp,
                        (dvoid **)&authp, (ub4) OCI_HTYPE_SESSION,
                        (size_t) 0, (dvoid **) 0);

  (void) OCIAttrSet((dvoid *) authp, (ub4) OCI_HTYPE_SESSION,
                 (dvoid *) "samp", (ub4)4,
                 (ub4) OCI_ATTR_USERNAME, errhp);

  (void) OCIAttrSet((dvoid *) authp, (ub4) OCI_HTYPE_SESSION,
                 (dvoid *) "password", (ub4) 4,
                 (ub4) OCI_ATTR_PASSWORD, errhp);

  /* Begin a User Session */
  checkerr(errhp, OCISessionBegin ( svchp,  errhp, authp, OCI_CRED_RDBMS,
                          (ub4) OCI_DEFAULT));

  (void) OCIAttrSet((dvoid *) svchp, (ub4) OCI_HTYPE_SVCCTX,
                   (dvoid *) authp, (ub4) 0,
                   (ub4) OCI_ATTR_SESSION, errhp);

  /* -----------------------User Logged In-------------------------------*/
  printf ("user logged in \n");

  /* allocate a statement handle */
  checkerr(errhp, OCIHandleAlloc( (dvoid *) envhp, (dvoid **) &stmthp,
           OCI_HTYPE_STMT, (size_t) 0, (dvoid **) 0));

  checkerr(errhp, OCIDescriptorAlloc((dvoid *)envhp, (dvoid **) &Lob_loc,
                                     (ub4) OCI_DTYPE_LOB,
                                     (size_t) 0, (dvoid **) 0));

  /* ------- subprogram called  here------------------------*/
  printf ("calling concat...\n");
  concat(ctx, str1, str1_i, str2, str2_i, ret_i, ret_l);

  return 0;
}

#endif
```

OCIExtProcRaiseExcp

This service routine raises a predefined exception, which must have a valid Oracle Database error number in the range 1..32,767. After doing any necessary cleanup, your external procedure must return immediately. (No values are assigned to `OUT` or `IN OUT` parameters.) The C prototype for this function follows:

```
int OCIExtProcRaiseExcp(
   OCIExtProcContext *with_context,
   size_t errnum);
```

The parameters `with_context` and `error_number` are the context pointer and Oracle Database error number. The return values `OCIEXTPROC_SUCCESS` and `OCIEXTPROC_ERROR` indicate success or failure.

In SQL*Plus, suppose you publish external procedure `plsTo_divide_proc`, as follows:

```
CREATE OR REPLACE PROCEDURE plsTo_divide_proc (
   dividend IN PLS_INTEGER,
   divisor  IN PLS_INTEGER,
   result   OUT FLOAT)
AS LANGUAGE C
   NAME "C_divide"
   LIBRARY MathLib
   WITH CONTEXT
   PARAMETERS (
      CONTEXT,
      dividend INT,
      divisor  INT,
      result   FLOAT);
```

When called, `C_divide` finds the quotient of two numbers. As this example shows, if the divisor is zero, `C_divide` uses `OCIExtProcRaiseExcp` to raise the predefined exception `ZERO_DIVIDE`:

```
void C_divide (ctx, dividend, divisor, result)
OCIExtProcContext *ctx;
int    dividend;
int    divisor;
float  *result;
{
  /* Check for zero divisor. */
  if (divisor == (int)0)
  {
    /* Raise exception ZERO_DIVIDE, which is Oracle Database error 1476. */
    if (OCIExtProcRaiseExcp(ctx, (int)1476) == OCIEXTPROC_SUCCESS)
    {
      return;
    }
    else
    {
      /* Incorrect parameters were passed. */
      assert(0);
    }
  }
  *result = (float)dividend / (float)divisor;
}
```

OCIExtProcRaiseExcpWithMsg

This service routine raises a user-defined exception and returns a user-defined error message. The C prototype for this function follows:

```
int OCIExtProcRaiseExcpWithMsg(
   OCIExtProcContext *with_context,
   size_t  error_number,
   text   *error_message,
   size_t  len);
```

The parameters `with_context`, `error_number`, and `error_message` are the context pointer, Oracle Database error number, and error message text. The parameter `len` stores the length of the error message. If the message is a null-terminated string, then `len` is zero. The return values `OCIEXTPROC_SUCCESS` and `OCIEXTPROC_ERROR` indicate success or failure.

In the previous example, you published external procedure `plsTo_divide_proc`. In this example, you use a different implementation. With this version, if the divisor is zero, then `C_divide` uses `OCIExtProcRaiseExcpWithMsg` to raise a user-defined exception:

```
void C_divide (ctx, dividend, divisor, result)
OCIExtProcContext *ctx;
int    dividend;
int    divisor;
float  *result;
  /* Check for zero divisor. */
  if (divisor == (int)0)
  {
    /* Raise a user-defined exception, which is Oracle Database error 20100,
       and return a null-terminated error message. */
    if (OCIExtProcRaiseExcpWithMsg(ctx, (int)20100,
          "divisor is zero", 0) == OCIEXTPROC_SUCCESS)
    {
      return;
    }
    else
    {
      /*  Incorrect parameters were passed. */
      assert(0);
    }
  }
  *result = dividend / divisor;

}
```

Doing Callbacks with External C Procedures

To enable callbacks, use the function OCIExtProcGetEnv.

Topics:

- OCIExtProcGetEnv
- Object Support for OCI Callbacks
- Restrictions on Callbacks
- Debugging External Procedures
- Example: Calling an External Procedure
- Global Variables in External C Procedures
- Static Variables in External C Procedures
- Restrictions on External C Procedures

OCIExtProcGetEnv

This service routine enables OCI callbacks to the database during an external procedure call. The environment handles obtained by using this function reuse the existing connection to go back to the database. If you must establish a new connection to the database, you cannot use these handles; instead, you must create your own.

The C prototype for this function follows:

```
sword OCIExtProcGetEnv ( OCIExtProcContext *with_context,
   OCIEnv envh,
```

```
   OCISvcCtx svch,
   OCIError errh )
```

The parameter `with_context` is the context pointer, and the parameters `envh`, `svch`, and `errh` are the OCI environment, service, and error handles, respectively. The return values `OCIEXTPROC_SUCCESS` and `OCIEXTPROC_ERROR` indicate success or failure.

Both external C procedures and Java class methods can call-back to the database to do SQL operations. For a working example, see "Example: Calling an External Procedure" on page 14-39 .

> **Note:** Callbacks are not necessarily a same-session phenomenon; you might run an SQL statement in a different session through `OCIlogon`.

An external C procedure running on Oracle Database can call a service routine to obtain OCI environment and service handles. With the OCI, you can use callbacks to run SQL statements and PL/SQL subprograms, fetch data, and manipulate LOBs. Callbacks and external procedures operate in the same user session and transaction context, and so have the same user privileges.

In SQL*Plus, suppose you run this script:

```
CREATE TABLE Emp_tab (empno NUMBER(10))

CREATE PROCEDURE plsToC_insertIntoEmpTab_proc (
   empno PLS_INTEGER)
AS LANGUAGE C
   NAME "C_insertEmpTab"
   LIBRARY insert_lib
   WITH CONTEXT
   PARAMETERS (
      CONTEXT,
      empno LONG);
```

Later, you might call service routine `OCIExtProcGetEnv` from external procedure `plsToC_insertIntoEmpTab_proc`, as follows:

```
#include <stdio.h>
#include <stdlib.h>
#include <oratypes.h>
#include <oci.h>   /* includes ociextp.h */
...
void C_insertIntoEmpTab (ctx, empno)
OCIExtProcContext *ctx;
long empno;
{
  OCIEnv    *envhp;
  OCISvcCtx *svchp;
  OCIError  *errhp;
  int        err;
  ...
  err = OCIExtProcGetEnv(ctx, &envhp, &svchp, &errhp);
  ...
}
```

If you do not use callbacks, you need not include `oci.h`; instead, just include `ociextp.h`.

Object Support for OCI Callbacks

To run object-related callbacks from your external procedures, the OCI environment in the `extproc` agent is fully initialized in object mode. You retrieve handles to this environment with the `OCIExtProcGetEnv` procedure.

The object run-time environment lets you use static and dynamic object support provided by OCI. To use static support, use the OTT to generate C structs for the appropriate object types, and then use conventional C code to access the object attributes.

For those objects whose types are unknown at external procedure creation time, an alternative, dynamic, way of accessing objects is first to call `OCIDescribeAny` to obtain attribute and method information about the type. Then, `OCIObjectGetAttr` and `OCIObjectSetAttr` can be called to retrieve and set attribute values.

Because the current external procedure model is stateless, `OCIExtProcGetEnv` must be called in every external procedure that wants to run callbacks, or call `OCIExtProc`. service routines. After every external procedure call, the callback mechanism is cleaned up and all OCI handles are freed.

Restrictions on Callbacks

With callbacks, this SQL statements and OCI subprograms are not supported:

- Transaction control statements such as `COMMIT`
- Data definition statements such as `CREATE`
- These object-oriented OCI subprograms:

```
OCIObjectNew
OCIObjectPin
OCIObjectUnpin
OCIObjectPinCountReset
OCIObjectLock
OCIObjectMarkUpdate
OCIObjectUnmark
OCIObjectUnmarkByRef
OCIObjectAlwaysLatest
OCIObjectNotAlwaysLatest
OCIObjectMarkDeleteByRef
OCIObjectMarkDelete
OCIObjectFlush
OCIObjectFlushRefresh
OCIObjectGetTypeRef
OCIObjectGetObjectRef
OCIObjectExists
OCIObjectIsLocked
OCIObjectIsDirtied
OCIObjectIsLoaded
OCIObjectRefresh
OCIObjectPinTable
OCIObjectArrayPin
OCICacheFlush,
OCICacheFlushRefresh,
OCICacheRefresh
OCICacheUnpin
OCICacheFree
OCICacheUnmark
OCICacheGetObjects
OCICacheRegister
```

- Polling-mode OCI subprograms such as `OCIGetPieceInfo`
- These OCI subprograms:

```
OCIEnvInit
OCIInitialize
OCIPasswordChange
OCIServerAttach
OCIServerDetach
OCISessionBegin
OCISessionEnd
OCISvcCtxToLda
OCITransCommit
OCITransDetach
OCITransRollback
OCITransStart
```

Also, with OCI subprogram `OCIHandleAlloc`, these handle types are not supported:

```
OCI_HTYPE_SERVER
OCI_HTYPE_SESSION
OCI_HTYPE_SVCCTX
OCI_HTYPE_TRANS
```

Debugging External Procedures

Usually, when an external procedure fails, its prototype is faulty. In other words, the prototype does not match the one generated internally by PL/SQL. This can happen if you specify an incompatible C data type. For example, to pass an `OUT` parameter of type `REAL`, you must specify `float *`. Specifying `float`, `double *`, or any other C data type results in a mismatch.

In such cases, you might get:

```
lost RPC connection to external routine agent
```

This error, which means that `extproc` terminated abnormally because the external procedure caused a core dump. To avoid errors when declaring C prototype parameters, see the preceding tables.

To help you debug external procedures, PL/SQL provides the utility package `DEBUG_EXTPROC`. To install the package, run the script `dbgextp.sql`, which you can find in the PL/SQL demo directory. (For the location of the directory, see your Oracle Database Installation or User's Guide.)

To use the package, follow the instructions in `dbgextp.sql`. Your Oracle Database account must have `EXECUTE` privileges on the package and `CREATE LIBRARY` privileges.

Note: `DEBUG_EXTPROC` works only on platforms with debuggers that can attach to a running process.

Example: Calling an External Procedure

Also in the PL/SQL demo directory is the script `extproc.sql`, which demonstrates the calling of an external procedure. The companion file `extproc.c` contains the C source code for the external procedure.

To run the demo, follow the instructions in `extproc.sql`. You must use the `SCOTT` account, which must have `CREATE LIBRARY` privileges.

Global Variables in External C Procedures

A global variable is declared outside of a function, and its value is shared by all functions of a program. Therefore, in external procedures, all functions in a DLL share the value of the global variable. Global variables are also used to store data that is intended to persist beyond the lifetime of a function. However, Oracle discourages the use of global variables for two reasons:

- Threading

  In the nonthreaded configuration of the agent process, one function is active at a time. For the multithreaded `extproc` agent, multiple functions can be active at the same time, and they might try to access the global variable concurrently, with unsuccessful results.

- DLL caching

  Suppose that function `func1` tries to pass data to function `func2` by storing the data in a global variable. After `func1` completes, the DLL cache might be unloaded, causing all global variables to lose their values. Then, when `func2` runs, the DLL is reloaded, and all global variables are initialized to 0.

Static Variables in External C Procedures

There are two types of static variables: external and internal. An external static variable is a special case of a global variable, so its usage is discouraged. Internal static variables are local to a particular function, but remain in existence rather than coming and going each time the function is activated. Therefore, they provide private, permanent storage within a single function. These variables are used to pass on data to subsequent calls to the same function. But, because of the DLL caching feature mentioned previously, the DLL might be unloaded and reloaded between calls, which means that the internal static variable loses its value.

> **See Also:** Template `makefile` in the RDBMS subdirectory `/public` for help creating a dynamic link library

When calling external procedures:

- Never write to `IN` parameters or overflow the capacity of `OUT` parameters. (PL/SQL does no run time checks for these error conditions.)
- Never read an `OUT` parameter or a function result.
- Always assign a value to `IN OUT` and `OUT` parameters and to function results. Otherwise, your external procedure will not return successfully.
- If you include the `WITH CONTEXT` and `PARAMETERS` clauses, then you must specify the parameter `CONTEXT`, which shows the position of the context pointer in the parameter list.
- If you include the `PARAMETERS` clause, and if the external procedure is a function, then you must specify the parameter `RETURN` in the last position.
- For every formal parameter, there must be a corresponding parameter in the `PARAMETERS` clause. Also, ensure that the data types of parameters in the `PARAMETERS` clause are compatible with those in the C prototype, because no implicit conversions are done.

- With a parameter of type `RAW` or `LONG RAW`, you must use the property `LENGTH`. Also, if that parameter is `IN OUT` or `OUT` and null, then you must set the length of the corresponding C parameter to zero.

Restrictions on External C Procedures

These restrictions apply to external procedures:

- This feature is available only on platforms that support DLLs.
- Only C procedures and procedures callable from C code are supported.
- External procedure callouts combined with distributed transactions is not supported.
- You cannot pass PL/SQL cursor variables or records to an external procedure. For records, use instances of object types instead.
- In the `LIBRARY` subclause, you cannot use a database link to specify a remote library.
- The maximum number of parameters that you can pass to a external procedure is 128. However, if you pass float or double parameters by value, then the maximum is less than 128. How much less depends on the number of such parameters and your operating system. To get a rough estimate, count each float or double passed by value as two parameters.

15

Developing Applications with Oracle XA

This chapter explains how to use the Oracle XA library. Typically, you use this library in applications that work with transaction monitors. The XA features are most useful in applications in which transactions interact with multiple databases.

Topics:

- X/Open Distributed Transaction Processing (DTP)
- Oracle XA Library Subprograms
- Developing and Installing XA Applications
- Troubleshooting XA Applications
- Oracle XA Issues and Restrictions

See Also:

- *X/Open CAE Specification - Distributed Transaction Processing: The XA Specification*, X/Open Document Number XO/CAE/91/300, for an overview of XA, including basic architecture. Access at `http://www.opengroup.org/pubs/catalog/c193.htm`.
- *Oracle Call Interface Programmer's Guide* for background and reference information about the Oracle XA library
- The Oracle Database platform-specific documentation for information about library linking filenames
- README for changes, bugs, and restrictions in the Oracle XA library for your platform

X/Open Distributed Transaction Processing (DTP)

The X/Open Distributed Transaction Processing (DTP) architecture defines a standard architecture or interface that enables multiple application programs (APs) to share resources provided by multiple, and possibly different, resource managers (RMs). It coordinates the work between APs and RMs into global transactions.

The Oracle XA library conforms to the X/Open software architecture's XA interface specification. The Oracle XA library is an external interface that enables a client-side transaction manager (TM) that is not an Oracle client-side TM to coordinate global transactions, thereby allowing inclusion of database RMs that are not Oracle Database RMs in distributed transactions. For example, a client application can manage an Oracle Database transaction and a transaction in an NTFS file system as a single, global transaction.

Figure 15–1 illustrates a possible X/Open DTP model.

Figure 15–1 Possible DTP Model

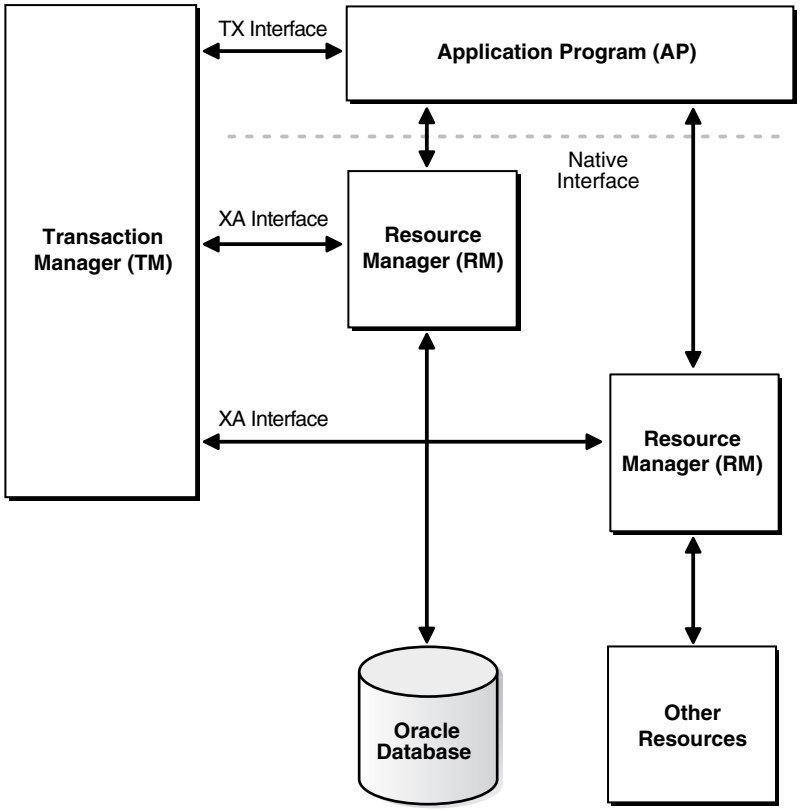

Topics:

- DTP Terminology
- Required Public Information

DTP Terminology

- Resource Manager (RM)
- Distributed Transaction
- Branch
- Transaction Manager (TM)
- Transaction Processing Monitor (TPM)
- Two-Phase Commit Protocol
- Application Program (AP)
- TX Interface
- Tight and Loose Coupling
- Dynamic and Static Registration

Resource Manager (RM)

A resource manager controls a shared, recoverable resource that can be returned to a consistent state after a failure. Examples are relational databases, transactional queues, and transactional file systems. Oracle Database is an RM and uses its online redo log and undo segments to return to a consistent state after a failure.

Distributed Transaction

A distributed transaction, also called a **global transaction**, is a client transaction that involves updates to multiple distributed resources and requires "all-or-none" semantics across distributed RMs.

Branch

A branch is a unit of work contained within one RM. Multiple branches comprise a global transaction. For Oracle Database, each branch maps to a local transaction inside the database server.

Transaction Manager (TM)

A transaction manager provides an API for specifying the boundaries of the transaction and manages commit and recovery. The TM implements a two-phase commit engine to provide "all-or-none" semantics across distributed RMs.

An **external TM** is a middle-tier component that resides outside Oracle Database. Normally, the database is its own internal TM. Using a standards-based TM enables Oracle Database to cooperate with other heterogeneous RMs in a single transaction.

Transaction Processing Monitor (TPM)

A TM is usually provided by a transaction processing monitor (TPM), such as:

- Oracle Tuxedo
- IBM Transarc Encina
- IBM CICS

A TPM coordinates the flow of transaction requests between the client processes that issue requests and the back-end servers that process them. Basically, a TPM coordinates transactions that require the services of several different types of back-end processes, such as application servers and RMs distributed over a network.

The TPM synchronizes any commits or rollbacks required to complete a distributed transaction. The TM portion of the TPM is responsible for controlling when distributed commits and rollbacks take place. Thus, if a distributed application program takes advantage of a TPM, then the TM portion of the TPM is responsible for controlling the two-phase commit protocol. The RMs enable the TMs to perform this task.

Because the TM controls distributed commits or rollbacks, it must communicate directly with Oracle Database (or any other RM) through the XA interface. It uses Oracle XA library subprograms, which are described in "Oracle XA Library Subprograms" on page 15-5, to tell Oracle Database how to process the transaction, based on its knowledge of all RMs in the transaction.

Two-Phase Commit Protocol

The Oracle XA library interface follows the two-phase commit protocol. The sequence of events is as follows:

1. In the prepare phase, the TM asks each RM to guarantee that it can commit any part of the transaction. If this is possible, then the RM records its prepared state

and replies affirmatively to the TM. If it is not possible, then the RM might roll back any work, reply negatively to the TM, and forget about the transaction. The protocol allows the application, or any RM, to roll back the transaction unilaterally until the prepare phase completes.

2. In phase two, the TM records the commit decision and issues a commit or rollback to all RMs participating in the transaction. TM can issue a commit for an RM only if all RMs have replied affirmatively to phase one.

Application Program (AP)

An application program defines transaction boundaries and specifies actions that constitute a transaction. For example, an AP can be a precompiler or Oracle Call Interface (OCI) program. The AP operates on the RM resource through its native interface, for example, SQL.

TX Interface

An application program starts and completes all transaction control operations through the TM through an interface called **TX**. The AP does not directly use the XA interface. APs are not aware of branches that fork in the middle-tier: application threads do not explicitly join, leave, suspend, and resume branch work, instead the TM portion of the transaction processing monitor manages the branches of a global transaction for APs. Ultimately, APs call the TM to commit all-or-none.

Note: The naming conventions for the TX interface and associated subprograms are vendor-specific. For example, the `tx_open` call might be referred to as `tp_open` on your system. In some cases, the calls might be implicit, for example, at the entry to a transactional RPC. See the documentation supplied with the transaction processing monitor for details.

Tight and Loose Coupling

Application threads are **tightly coupled** if the RM considers them as a single entity for all isolation semantic purposes. Tightly coupled branches must see changes in each other. Furthermore, an external client must either see all changes of a tightly coupled set or none of the changes. If application threads are not tightly coupled, then they are **loosely coupled**.

Dynamic and Static Registration

Oracle Database supports both dynamic and static registration. In **dynamic registration**, the RM runs an application callback before starting any work. In **static registration**, you must call `xa_start` for each RM before starting any work, even if some RMs are not involved.

Required Public Information

As a resource manager, Oracle Database must publish the information described in Table 15–1.

Table 15–1 Required XA Features Published by Oracle Database

XA Feature	Oracle Database Details
`xa_switch_t` structures	The Oracle Database `xa_switch_t` structure name is `xaosw` for static registration and `xaoswd` for dynamic registration. These structures contain entry points and other information for the resource manager.
`xa_switch_t` resource manager	The Oracle Database resource manager name within the `xa_switch_t` structure is `Oracle_XA`.
Close string	The close string used by `xa_close` is ignored and can be null.
Open string	For the description of the format of the open string that `xa_open` uses, see "Defining the xa_open String" on page 15-8.
Libraries	Libraries needed to link applications using Oracle XA have platform-specific names. The procedure is similar to linking an ordinary precompiler or OCI program except that you might have to link any TPM-specific libraries. If you are not using `sqllib`, then link with `$ORACLE_HOME/rdbms/lib/xaonsl.o` or `$ORACLE_HOME/rdbms/lib32/xaonsl.o` (for 32 bit application on 64 bit platforms).
Requirements	None. The functionality to support XA is part of both Standard Edition and Enterprise Edition.

Oracle XA Library Subprograms

The Oracle XA library subprograms enable a TM to tell Oracle Database how to process transactions. Generally, the TM must open the resource by using `xa_open`. Typically, the opening of the resource results from the AP call to `tx_open`. Some TMs might call `xa_open` implicitly when the application begins.

Similarly, there is a close (using `xa_close`) that occurs when the application is finished with the resource. The close might occur when the AP calls `tx_close` or when the application terminates.

The TM instructs the RMs to perform several other tasks, which include:

- Starting a transaction and associating it with an ID
- Rolling back a transaction
- Preparing and committing a transaction

Topics:

- Oracle XA Library Subprograms
- Oracle XA Interface Extensions

Oracle XA Library Subprograms

XA Library subprograms are described in Table 15–2.

Table 15–2 XA Library Subprograms

XA Subprogram	Description
`xa_open`	Connects to the RM.
`xa_close`	Disconnects from the RM.

Table 15–2 (Cont.) XA Library Subprograms

XA Subprogram	Description
`xa_start`	Starts a transaction and associates it with the given transaction ID (XID), or associates the process with an existing transaction.
`xa_end`	Disassociates the process from the given XID.
`xa_rollback`	Rolls back the transaction associated with the given XID.
`xa_prepare`	Prepares the transaction associated with the given XID. This is the first phase of the two-phase commit protocol.
`xa_commit`	Commits the transaction associated with the given XID. This is the second phase of the two-phase commit protocol.
`xa_recover`	Retrieves a list of prepared, heuristically committed, or heuristically rolled back transactions.
`xa_forget`	Forgets the heuristically completed transaction associated with the given XID.

In general, the AP need not worry about the subprograms in Table 15–2 except to understand the role played by the `xa_open` string.

Oracle XA Interface Extensions

Oracle Database's XA interface includes some additional functions, which are described in Table 15–3.

Table 15–3 Oracle XA Interface Extensions

Function	Description
`OCISvcCtx *xaoSvcCtx(text *dbname)`	Returns the OCI service handle for a given XA connection. The `dbname` parameter must be the same as the `DB` parameter passed in the `xa_open` string. OCI applications can use this routing instead of the `sqlld2` calls to obtain the connection handle. Hence, OCI applications need not link with the `sqllib` library. The service handle can be converted to the Version 7 OCI logon data area (LDA) by using `OCISvcCtxToLda` [Version 8 OCI]. Client applications must remember to convert the Version 7 LDA to a service handle by using `OCILdaToSvcCtx` after completing the OCI calls.
`OCIEnv *xaoEnv(text *dbname)`	Returns the OCI environment handle for a given XA connection. The `dbname` parameter must be the same as the `DB` parameter passed in the `xa_open` string.
`int xaosterr(OCISvcCtx *SvcCtx,sb4 error)`	Converts an Oracle Database error code to an XA error code (only applicable to dynamic registration). The first parameter is the service handle used to run the work in the database. The second parameter is the error code that was returned from Oracle Database. Use this function to determine if the error returned from an OCI statement was caused because the `xa_start` failed. The function returns XA_OK if the error was not generated by the XA module or a valid XA error if the error was generated by the XA module.

Developing and Installing XA Applications

This section explains how to develop and install Oracle XA applications:

- DBA or System Administrator Responsibilities
- Application Developer Responsibilities
- Defining the xa_open String
- Developing and Installing XA Applications
- Managing Transaction Control with Oracle XA
- Migrating Precompiler or OCI Applications to TPM Applications
- Managing Oracle XA Library Thread Safety
- Using the DBMS_XA Package

DBA or System Administrator Responsibilities

The responsibilities of the DBA or system administrator are as follows:

1. Define the open string, with help from the application developer. For details, see "Defining the xa_open String" on page 15-8.

2. Ensure that the static data dictionary view `DBA_PENDING_TRANSACTIONS` exists and grant the `SELECT` privilege to the view for all Oracle users specified in the `xa_open` string.

   Grant `FORCE TRANSACTION` privilege to the Oracle user who might commit or roll back pending (in-doubt) transactions that he or she created, using the command `COMMIT FORCE` *`local_tran_id`* or `ROLLBACK FORCE` *`local_tran_id`*.

   Grant `FORCE ANY TRANSACTION` privilege to the Oracle user who might commit or roll back XA transactions created by other users. For example, if user A might commit or roll back a transaction that was created by user B, user A must have `FORCE ANY TRANSACTION` privilege.

   In Oracle Database version 7 client applications, all Oracle Database accounts used by Oracle XA library must have the `SELECT` privilege on the dynamic performance view `V$XATRANS$`. This view must have been created during the XA library installation. If necessary, you can manually create the view by running the SQL script `xaview.sql` as Oracle Database user `SYS`.

   > **See Also:** Your Oracle Database platform-specific documentation for the location of the `catxpend.sql` script

3. Using the open string information, install the RM into the TPM configuration. Follow the TPM vendor instructions.

   The DBA or system administrator must be aware that a TPM system starts the process that connects to Oracle Database. See your TPM documentation to determine what environment exists for the process and what user ID it must have. Ensure that correct values are set for `$ORACLE_HOME` and `$ORACLE_SID` in this environment.

4. Grant the user ID write permission to the directory in which the system is to write the XA trace file.

   > **See Also:** "Defining the xa_open String" on page 15-8 for information about how to specify an Oracle System Identifier (SID) or a trace directory that is different from the defaults

5. Start the relevant database instances to bring Oracle XA applications on-line. Perform this task before starting any TPM servers.

Application Developer Responsibilities

The responsibilities of the application developer are as follows:

1. Define the open string with help from the DBA or system administrator, as explained in "Defining the xa_open String" on page 15-8.
2. Develop the applications.

   Observe special restrictions on transaction-oriented SQL statements for precompilers.

   See Also: "Developing and Installing XA Applications" on page 15-6

3. Link the application according to TPM vendor instructions.

Defining the xa_open String

The open string is used by the transaction monitor to open the database. The maximum number of characters in an open string is 256.

Topics:

- Syntax of the xa_open String
- Required Fields for the xa_open String
- Optional Fields for the xa_open String

Syntax of the xa_open String

You can define an open string with the syntax shown in Example 15–1.

Example 15–1 xa_open String

```
ORACLE_XA{+required_fields...} [+optional_fields...]
```

These strings shows sample parameter settings:

```
ORACLE_XA+DB=MANAGERS+SqlNet=SID1+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/xalog
ORACLE_XA+DB=PAYROLL+SqlNet=SID2+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/xalog
ORACLE_XA+SqlNet=SID3+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/xalog
```

These topics describe valid parameters for the *required_fields* and *optional_fields* placeholders:

- Required Fields for the xa_open String
- Optional Fields for the xa_open String

Note:

- You can enter the required fields and optional fields in any order when constructing the open string.
- All field names are case insensitive. Whether their values are case-sensitive depends on the platform.
- There is no way to use the plus character (+) as part of the actual information string.

Required Fields for the xa_open String

The *required_fields* placeholder in Example 15–1 refers to any of the name-value pairs described in Table 15–4.

Table 15–4 Required Fields of xa_open string

Syntax Element	Description
`Acc=P//`	Specifies that no explicit user or password information is provided and that the operating system authentication form is used. For more information see *Oracle Database Administrator's Guide*.
`Acc=P/`*`user/password`*	Specifies the user name and password for a valid Oracle Database account. As described in "DBA or System Administrator Responsibilities" on page 15-7, ensure that HR has the `SELECT` privilege on the `DBA_PENDING_TRANSACTIONS` table.
`SesTm=`*`session_time_limit`*	Specifies the maximum number of seconds allowed in a transaction between one service and the next, or between a service and the commit or rollback of the transaction, before the system terminates the transaction. For example, `SesTM=15` indicates that the session idle time limit is 15 seconds. For example, if the TPM uses remote subprogram calls between the client and the servers, then `SesTM` applies to the time between the completion of one RPC and the initiation of the next RPC, or the `tx_commit`, or the `tx_rollback`. The value of `0` indicates no limit. Entering a value of `0` is strongly discouraged. It might tie up resources for a long time if something goes wrong. Also, if a child process has `SesTM=0`, then the `SesTM` setting is not effective after the parent process is terminated.

Optional Fields for the xa_open String

The *optional_fields* placeholder in Example 15–1 refers to any of the name-value pairs described in Table 15–5.

Table 15–5 Optional Fields in the xa_open String

Syntax Element	Description
`NoLocal= true | false`	Specifies whether local transactions are allowed. The default value is `false`. If the application must disallow local transactions, then set the value to `true`.

Table 15–5 (Cont.) Optional Fields in the xa_open String

Syntax Element	Description
DB=*db_name*	Specifies the name used by Oracle Database precompilers to identify the database. For example, DB=payroll specifies that the database name is payroll and that the application server program uses that name in AT clauses. Application programs that use only the default database for the Oracle Database precompiler (that is, they do not use the AT clause in their SQL statements) must omit the DB=*db_name* clause in the open string. Applications that use explicitly named databases must indicate that database name in their DB=*db_name* field. Oracle Database Version 7 OCI programs must call the sqlld2 function to obtain the correct context for logon data area (Lda_Def), which is the equivalent of an OCI service context. Version 8 and higher OCI programs must call the xaoSvcCtx function to get the OCISvcCtx service context. The *db_name* is not the SID and is not used to locate the database to be opened. Rather, it correlates the database opened by this open string with the name used in the application program to run SQL statements. The SID is set from either the environment variable ORACLE_SID of the TPM application server or the SID given in the Oracle Net clause in the open string. The Oracle Net clause is described later in this section.Some TPM vendors provide a way to name a group of servers that use the same open string. You might find it convenient to choose the same name both for that purpose and for db_name.
LogDir=*log_dir*	Specifies the path name on the local system where the Oracle XA library error and tracing information is tomust be logged. The default is $ORACLE_HOME/rdbms/log if ORACLE_HOME is set; otherwise, it specifies the current directory. For example, LogDir=/xa_trace indicates that the logging information is located under the /xa_trace directory. Ensure that the directory exists and the application server can write to it.
Objects= true | false	Specifies whether the application is initialized in object mode. The default value is false. If the application must use certain API calls that require object mode, such as OCIRawAssignBytes, then set the value to true.
MaxCur=*maximum_#_of_open_cursors*	Specifies the number of cursors to be allocated when the database is opened. It serves the same purpose as the precompiler option maxopencursors. For example, MaxCur=5 indicates that the precompiler tries to keep five open cursors cached. This parameter overrides the precompiler option maxopencursors that you might have specified in your source code or at compile time.
SqlNet=*db_link*	Specifies the Oracle Net database link to use to log on to the system. This string must be an entry in tnsnames.ora. For example, the string SqlNet=inst1_disp might connect to a shared server at instance 1 if so defined in tnsnames.ora. You can use the SqlNet parameter to specify the ORACLE_SID in cases where you cannot control the server environment variable. You must also use it when the server must access multiple Oracle Database instances. To use the Oracle Net string without actually accessing a remote database, use the Pipe driver. For example, specify SqlNet=localsid1, where localsid1 is an alias defined in the tnsnames.ora file.

Table 15–5 (Cont.) Optional Fields in the xa_open String

Syntax Element	Description
`Loose_Coupling=true | false`	Specifies whether locks are shared. Oracle Database transaction branches within the same global transaction can be coupled tightly or loosely. If branches are loosely coupled, then they do not share locks. Set the value to `true` for loosely coupled branches. If branches are tightly coupled, then they share locks. Set the value to `false` for tightly coupled branches. The default value is `false`.
`SesWt=`*`session_wait_limit`*	Specifies the number of seconds Oracle Database waits for a transaction branch that is being used by another session before `XA_RETRY` is returned. The default value is 60 seconds.
`Threads=true | false`	Specifies whether the application is multithreaded. The default value is `false`. If the application is multithreaded, then the setting is `true`.

Using Oracle XA with Precompilers

When used in an Oracle XA application, cursors are valid only for the duration of the transaction. Explicit cursors must be opened after the transaction begins, and closed before the commit or rollback.

You have these options when interfacing with precompilers:

- Using Precompilers with the Default Database
- Using Precompilers with a Named Database

The examples in this topic use the precompiler Pro*C/C++.

Using Precompilers with the Default Database

To interface to a precompiler with the default database, ensure that the `DB=`*`db_name`* field used in the open string is not present. The absence of this field indicates the default connection. Only one default connection is allowed for each process.

This is an example of an open string identifying a default Pro*C/C++ connection:

```
ORACLE_XA+SqlNet=maildb+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/logs
```

The `DB=`*`db_name`* is absent, indicating an empty database ID string.

The syntax of a SQL statement is:

```
EXEC SQL UPDATE Emp_tab SET Sal = Sal*1.5;
```

Using Precompilers with a Named Database

To interface to a precompiler with a named database, include the `DB=`*`db_name`* field in the open string. Any database you refer to must reference the same *`db_name`* you specified in the corresponding open string.

An application might include the default database and one or more named databases. For example, suppose you want to update an employee's salary in one database, his department number (`DEPTNO`) in another, and his manager in a third database. Configure the open strings in the transaction manager as shown in Example 15–2.

Example 15–2 Sample Open String Configuration

```
ORACLE_XA+DB=MANAGERS+SqlNet=SID1+ACC=P/username/password
```

```
  +SesTM=10+LogDir=/usr/local/xalog
ORACLE_XA+DB=PAYROLL+SqlNet=SID2+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/xalog
ORACLE_XA+SqlNet=SID3+ACC=P/username/password
  +SesTM=10+LogDir=/usr/local/xalog
```

There is no `DB=db_name` field in the last open string in Example 15–2.

In the application server program, enter declarations such as:

```
EXEC SQL DECLARE PAYROLL DATABASE;
EXEC SQL DECLARE MANAGERS DATABASE;
```

Again, the default connection (corresponding to the third open string that does not contain the `DB` field) needs no declaration.

When doing the update, enter statements similar to these:

```
EXEC SQL AT PAYROLL UPDATE Emp_Tab SET Sal=4500 WHERE Empno=7788;
EXEC SQL AT MANAGERS UPDATE Emp_Tab SET Mgr=7566 WHERE Empno=7788;
EXEC SQL UPDATE Emp_Tab SET Deptno=30 WHERE Empno=7788;
```

There is no `AT` clause in the last statement because it is referring to the default database.

In Oracle Database precompilers release 1.5.3 or later, you can use a character host variable in the `AT` clause, as this example shows:

```
EXEC SQL BEGIN DECLARE SECTION;
  DB_NAME1 CHARACTER(10);
  DB_NAME2 CHARACTER(10);
EXEC SQL END DECLARE SECTION;
    ...
SET DB_NAME1 = 'PAYROLL'
SET DB_NAME2 = 'MANAGERS'
    ...
EXEC SQL AT :DB_NAME1 UPDATE...
EXEC SQL AT :DB_NAME2 UPDATE...
```

Caution: Do not have XA applications create connections other than those created through `xa_open`. Work performed on non-XA connections is outside the global transaction and must be committed separately.

Using Oracle XA with OCI

Oracle Call Interface applications that use the Oracle XA library must not call `OCISessionBegin` to log on to the resource manager. Rather, the logon must be done through the TPM. The applications can run the function `xaoSvcCtx` to obtain the service context structure when they must access the resource manager.

In applications that must pass the environment handle to OCI functions, you can also call `xaoEnv` to find that handle.

Because an application server can have multiple concurrent open Oracle Database resource managers, it must call the function `xaoSvcCtx` with the correct arguments to obtain the correct service context.

See Also: *Oracle Call Interface Programmer's Guide*

Managing Transaction Control with Oracle XA

When you use the XA library, transactions are not controlled by the SQL statements that commit or roll back transactions. Rather, they are controlled by an API accepted by the TM that starts and stops transactions. You call the API that is provided by the transaction manager, including the TX interface listed in Table 15–6, but not the XA Library Subprograms listed in Table 15–2.

The TMs typically control the transactions through the XA interface. This interface includes the functions described in Table 15–2.

Table 15–6 TX Interface Functions

TX Function	Description
`tx_open`	Logs into the resource manager(s)
`tx_close`	Logs out of the resource manager(s)
`tx_begin`	Starts a transaction
`tx_commit`	Commits a transaction
`tx_rollback`	Rolls back the transaction

Most TPM applications use a client/server architecture in which an application client requests services and an application server provides them. The examples shown in "Examples of Precompiler Applications" on page 15-13 use such a client/server model. A service is a logical unit of work that, for Oracle Database as the resource manager, comprises a set of SQL statements that perform a related unit of work.

For example, when a service named "credit" receives an account number and the amount to be credited, it runs SQL statements to update information in certain tables in the database. In addition, a service might request other services. For example, a "transfer fund" service might request services from a "credit" and "debit" service.

Typically, application clients request services from the application servers to perform tasks within a transaction. For some TPM systems, however, the application client itself can offer its own local services. As shown in "Examples of Precompiler Applications" on page 15-13, you can encode transaction control statements within either the client or the server.

To have multiple processes participating in the same transaction, the TPM provides a communication API that enables transaction information to flow between the participating processes. Examples of communications APIs include RPC, pseudo-RPC functions, and send/receive functions.

Because the leading vendors support different communication functions, these examples use the communication pseudo-function `tpm_service` to generalize the communications API.

X/Open includes several alternative methods for providing communication functions in their preliminary specification. At least one of these alternatives is supported by each of the leading TPM vendors.

Examples of Precompiler Applications

These examples illustrate precompiler applications. Assume that the application server has logged onto the RMs system, in a TPM-specific manner. Example 15–3 shows a transaction started by an application server.

Example 15–3 Transaction Started by an Application Server

```
/***** Client: *****/
tpm_service("ServiceName");              /*Request Service*/

/***** Server: *****/
ServiceName()
{
  <get service specific data>
  tx_begin();                             /* Begin transaction boundary */
  EXEC SQL UPDATE ...;

  /* This application server temporarily becomes */
  /* a client and requests another service. */

  tpm_service("AnotherService");
  tx_commit();                             /* Commit the transaction */
  <return service status back to the client>
}
```

Example 15–4 shows a transaction started by an application client.

Example 15–4 Transaction Started by an Application Client

```
/***** Client: *****/
tx_begin();                           /* Begin transaction boundary */
tpm_service("Service1");
tpm_service("Service2");
tx_commit();                          /* Commit the transaction */

/***** Server: *****/
Service1()
{
  <get service specific data>
  EXEC SQL UPDATE ...;
  <return service status back to the client>
}
Service2()
{
  <get service specific data>
  EXEC SQL UPDATE ...;
  ...
  <return service status back to client>
}
```

Migrating Precompiler or OCI Applications to TPM Applications

To migrate existing precompiler or OCI applications to a TPM application that uses the Oracle XA library, you must:

1. Reorganize the application into a framework of "services" so that application clients request services from application servers. Some TPMs require the application to use the `tx_open` and `tx_close` functions, whereas other TPMs do the logon and logoff implicitly.

   If you do not specify the `SqlNet` parameter in your open string, then the application uses the default Oracle Net driver. Thus, ensure that the application server is brought up with the `ORACLE_HOME` and `ORACLE_SID` environment variables properly defined. This is accomplished in a TPM-specific fashion. See your TPM vendor documentation for instructions on how to accomplish this.

2. Ensure that the application replaces the regular connect and disconnect statements. For example, replace the connect statements `EXEC SQL CONNECT` (for precompilers) or `OCISessionBegin`, `OCIServerAttach`, and `OCIEnvCreate` (for OCI) with `tx_open`. Replace the disconnect statements `EXEC SQL COMMIT/ROLLBACK WORK RELEASE` (for precompilers) or `OCISessionEnd/OCIServerDetach` (for OCI) with `tx_close`.

3. Ensure that the application replaces the regular commit or rollback statements for any global transactions and begins the transaction explicitly.

   For example, replace the `COMMIT/ROLLBACK` statements `EXEC SQL COMMIT/ROLLBACK WORK` (for precompilers), or `OCITransCommit/OCITransRollback` (for OCI) with `tx_commit/tx_rollback` and start the transaction by calling `tx_begin`.

   Note: The preceding is only true for global rather than local transactions. Commit or roll back local transactions with the Oracle API.

4. Ensure that the application resets the fetch state before ending a transaction. In general, use `release_cursor=no`. Use `release_cursor=yes` only when you are certain that a statement will run only once.

Table 15–7 lists the TPM functions that replace regular Oracle Database statements when migrating precompiler or OCI applications to TPM applications.

Table 15–7 TPM Replacement Statements

Regular Oracle Database Statements	TPM Functions
`CONNECT`*`user/password`*	`tx_open` (possibly implicit)
implicit start of transaction	`tx_begin`
`SQL`	Service that runs the SQL
`COMMIT`	`tx_commit`
`ROLLBACK`	`tx_rollback`
disconnect	`tx_close` (possibly implicit)

Managing Oracle XA Library Thread Safety

If you use a transaction monitor that supports threads, then the Oracle XA library enables you to write applications that are thread-safe. Nevertheless, keep certain issues in mind.

A **thread of control** (or thread) refers to the set of connections to resource managers. In an nonthreaded system, each process is considered a thread of control because each process has its own set of connections to RMs and maintains its own independent resource manager table. In a threaded system, each thread has an autonomous set of connections to RMs and each thread maintains a *private* RM table. This private table must be allocated for each thread and de-allocated when the thread terminates, even if the termination is abnormal.

Note: In Oracle Database, each thread that accesses the database must have its own connection.

Topics:

- Specifying Threading in the Open String
- Restrictions on Threading in Oracle XA

Specifying Threading in the Open String

The `xa_open` string provides the clause `Threads=`. You must specify this clause as `true` to enable the use of threads by the TM. The default is `false`. In most cases, the TM creates the threads; the application does not know when a thread is created. Therefore, it is advisable to allocate a service context on the stack within each service that is written for a TM application. Before doing any Oracle Database-related calls in that service, you must call the `xaoSvcCtx` function to retrieve the initialized OCI service context. You can then use this context for OCI calls within the service.

Restrictions on Threading in Oracle XA

These restrictions apply when using threads:

- Any Pro* or OCI code that runs as part of the application server process on the transaction monitor cannot be threaded unless the transaction monitor is explicitly told when each application thread is started. This is typically accomplished by using a special C compiler provided by the TM vendor.
- The Pro* statements `EXEC SQL ALLOCATE` and `EXEC SQL USE` are not supported. Therefore, when threading is enabled, you cannot use embedded SQL statements across non-XA connections.
- If one thread in a process connects to Oracle Database through XA, then all other threads in the process that connect to Oracle Database must also connect through XA. You cannot connect through `EXEC SQL CONNECT` in one thread and through `xa_open` in another thread.

Using the DBMS_XA Package

PL/SQL applications can use the Oracle XA library with the `DBMS_XA` package. For information about this package, see *Oracle Database PL/SQL Packages and Types Reference*.

In Example 15–5, one PL/SQL session starts a transaction but does not commit it, a second session resumes the transaction, and a third session commits the transaction. All three sessions are connected to the `HR` schema.

Example 15–5 Using the DBMS_XA Package

```
REM Session 1 starts a transaction and does some work.
DECLARE
  rc  PLS_INTEGER;
  oer PLS_INTEGER;
  xae EXCEPTION;
BEGIN
  rc  := DBMS_XA.XA_START(DBMS_XA_XID(123), DBMS_XA.TMNOFLAGS);

  IF rc!=DBMS_XA.XA_OK THEN
    oer := DBMS_XA.XA_GETLASTOER();
    DBMS_OUTPUT.PUT_LINE('ORA-' || oer || ' occurred, XA_START failed');
    RAISE xae;
  ELSE DBMS_OUTPUT.PUT_LINE('XA_START(new xid=123)     OK');
  END IF;
```

```
  UPDATE employees SET salary=salary*1.1 WHERE employee_id = 100;
  rc  := DBMS_XA.XA_END(DBMS_XA_XID(123), DBMS_XA.TMSUSPEND);

  IF rc!=DBMS_XA.XA_OK THEN
    oer := DBMS_XA.XA_GETLASTOER();
    DBMS_OUTPUT.PUT_LINE('ORA-' || oer || ' occurred, XA_END failed');
    RAISE xae;
  ELSE DBMS_OUTPUT.PUT_LINE('XA_END(suspend xid=123)   OK');
  END IF;

  EXCEPTION
    WHEN OTHERS THEN
      DBMS_OUTPUT.PUT_LINE
       ('XA error('||rc||') occurred, rolling back the transaction ...');
      rc := DBMS_XA.XA_END(DBMS_XA_XID(123), DBMS_XA.TMSUCCESS);
      rc := DBMS_XA.XA_ROLLBACK(DBMS_XA_XID(123));

      IF rc != DBMS_XA.XA_OK THEN
        oer := DBMS_XA.XA_GETLASTOER();
        DBMS_OUTPUT.PUT_LINE('XA-'||rc||', ORA-' || oer ||
         ' XA_ROLLBACK does not return XA_OK');
        raise_application_error(-20001, 'ORA-'||oer||
         ' error in rolling back a failed transaction');
      END IF;

      raise_application_error(-20002, 'ORA-'||oer||
       ' error in transaction processing, transaction rolled back');
END;
/
SHOW ERRORS
DISCONNECT

REM Session 2 resumes the transaction and does some work.
DECLARE
  rc  PLS_INTEGER;
  oer PLS_INTEGER;
  s   NUMBER;
  xae EXCEPTION;
BEGIN
  rc  := DBMS_XA.XA_START(DBMS_XA_XID(123), DBMS_XA.TMRESUME);

  IF rc!=DBMS_XA.XA_OK THEN
    oer := DBMS_XA.XA_GETLASTOER();
    DBMS_OUTPUT.PUT_LINE('ORA-' || oer || ' occurred, xa_start failed');
    RAISE xae;
  ELSE DBMS_OUTPUT.PUT_LINE('XA_START(resume xid=123)  OK');
  END IF;

  SELECT salary INTO s FROM employees WHERE employee_id = 100;
  DBMS_OUTPUT.PUT_LINE('employee_id = 100, salary = ' || s);
  rc  := DBMS_XA.XA_END(DBMS_XA_XID(123), DBMS_XA.TMSUCCESS);

  IF rc!=DBMS_XA.XA_OK THEN
    oer := DBMS_XA.XA_GETLASTOER();
    DBMS_OUTPUT.PUT_LINE('ORA-' || oer || ' occurred, XA_END failed');
    RAISE xae;
  ELSE DBMS_OUTPUT.PUT_LINE('XA_END(detach xid=123)    OK');
  END IF;

  EXCEPTION
```

```
    WHEN OTHERS THEN
      DBMS_OUTPUT.PUT_LINE
       ('XA error('||rc||') occurred, rolling back the transaction ...');
      rc := DBMS_XA.XA_END(DBMS_XA_XID(123), DBMS_XA.TMSUCCESS);
      rc := DBMS_XA.XA_ROLLBACK(DBMS_XA_XID(123));

      IF rc != DBMS_XA.XA_OK THEN
        oer := DBMS_XA.XA_GETLASTOER();
        DBMS_OUTPUT.PUT_LINE('XA-'||rc||', ORA-' || oer ||
         ' XA_ROLLBACK does not return XA_OK');
        raise_application_error(-20001, 'ORA-'||oer||
         ' error in rolling back a failed transaction');
      END IF;

      raise_application_error(-20002, 'ORA-'||oer||
       ' error in transaction processing, transaction rolled back');
END;
/
SHOW ERRORS
DISCONNECT

REM Session 3 commits the transaction.
DECLARE
  rc  PLS_INTEGER;
  oer PLS_INTEGER;
  xae EXCEPTION;
BEGIN
  rc  := DBMS_XA.XA_COMMIT(DBMS_XA_XID(123), TRUE);

  IF rc!=DBMS_XA.XA_OK THEN
    oer := DBMS_XA.XA_GETLASTOER();
    DBMS_OUTPUT.PUT_LINE('ORA-' || oer || ' occurred, XA_COMMIT failed');
    RAISE xae;
  ELSE DBMS_OUTPUT.PUT_LINE('XA_COMMIT(commit xid=123)  OK');
  END IF;

  EXCEPTION
    WHEN xae THEN
      DBMS_OUTPUT.PUT_LINE
       ('XA error('||rc||') occurred, rolling back the transaction ...');
      rc := DBMS_XA.XA_ROLLBACK(DBMS_XA_XID(123));

      IF rc != DBMS_XA.XA_OK THEN
        oer := DBMS_XA.XA_GETLASTOER();
        DBMS_OUTPUT.PUT_LINE('XA-'||rc||', ORA-' || oer ||
         ' XA_ROLLBACK does not return XA_OK');
        raise_application_error(-20001, 'ORA-'||oer||
         ' error in rolling back a failed transaction');
      END IF;

      raise_application_error(-20002, 'ORA-'||oer||
       ' error in transaction processing, transaction rolled back');
END;
/
SHOW ERRORS
DISCONNECT
QUIT
```

Troubleshooting XA Applications

Topics:

- Accessing Oracle XA Trace Files
- Managing In-Doubt or Pending Oracle XA Transactions
- Using SYS Account Tables to Monitor Oracle XA Transactions

Accessing Oracle XA Trace Files

The Oracle XA library logs any error and tracing information to its trace file. This information is useful in supplementing the XA error codes. For example, it can indicate whether an `xa_open` failure is caused by an incorrect open string, failure to find the Oracle Database instance, or a logon authorization failure.

The name of the trace file is `xa_`*`db_namedate`*`.trc`, where *`db_name`* is the database name specified in the open string field `DB=`*`db_name`*, and *`date`* is the date when the information is logged to the trace file. If you do not specify `DB=`*`db_name`* in the open string, then it automatically defaults to `NULL`.

For example, `xa_NULL06022005.trc` indicates a trace file that was created on June 2, 2005. Its `DB` field was not specified in the open string when the resource manager was opened. The filename `xa_Finance12152004.trc` indicates a trace file was created on December 15, 2004. Its `DB` field was specified as "Finance" in the open string when the resource manager was opened.

> **Note:** Multiple Oracle XA library resource managers with the same `DB` field and `LogDir` field in their open strings log all trace information that occurs on the same day to the same trace file.

Suppose that a trace file contains these contents:

```
1032.12345.2:  ORA-01017:  invalid username/password;  logon denied
1032.12345.2:  xaolgn:  XAER_INVAL;  logon denied
```

Table 15–8 explains the meaning of each element.

Table 15–8 Sample Trace File Contents

String	Description
`1032`	The time when the information is logged.
`12345`	The process ID (PID).
`2`	Resource manager ID
`xaolgn`	Name of module
`XAER_INVAL`	Error returned as specified in the XA standard
ORA-01017	Oracle Database information that was returned

Topics:

- xa_open String DbgFl
- Trace File Locations

xa_open String DbgFl

Normally, the XA trace file is opened only if an error is detected. The `xa_open` string `DbgFl` provides a tracing facility to record additional detail about the XA library. By default, its value is zero. You can set it to any combination of these values:

- `0x1`, which enables you to trace the entry and exit to each subprogram in the XA interface. This value can be useful in seeing exactly which XA calls the TP Monitor is making and which transaction identifier it is generating.
- `0x2`, which enables you to trace the entry to and exit from other nonpublic XA library programs. This is generally of use only to Oracle Database developers.
- `0x4`, which enables you to trace various other "interesting" calls made by the XA library, such as specific calls to the OCI. This is generally of use only to Oracle Database developers.

> **Note:** The flags are independent bits of an `ub4`, so to obtain printout from two or more flags, you must set a combined value of the flags.

Trace File Locations

The XA application determines a location for the trace file according to this algorithm:

1. The `LogDir` directory specified in the open string.
2. If you do not specify `LogDir` in the open string, then the Oracle XA application attempts to create the trace file in this directory (if the Oracle home is accessible):
   - `%ORACLE_HOME%\rdbms\trace` on Windows
   - `$ORACLE_HOME/rdbms/log` on Linux and UNIX
3. If the Oracle XA application cannot determine where the Oracle home is located, then the application creates the trace file in the current working directory.

Managing In-Doubt or Pending Oracle XA Transactions

In-doubt or pending transactions are transactions that were prepared but not committed to the database. In general, the TM provided by the TPM system resolves any failure and recovery of in-doubt or pending transactions. The DBA might have to override an in-doubt transaction if these situations occur:

- It is locking data that is required by other transactions.
- It is not resolved in a reasonable amount of time.

See the TPM documentation for more information about overriding in-doubt transactions in such circumstances and about how to decide whether to commit or roll back the in-doubt transaction.

Using SYS Account Tables to Monitor Oracle XA Transactions

These views under the Oracle Database `SYS` account contain transactions generated by regular Oracle Database applications and Oracle XA applications:

- `DBA_PENDING_TRANSACTIONS`
- `V$GLOBAL_TRANSACTION`
- `DBA_2PC_PENDING`
- `DBA_2PC_NEIGHBORS`

For transactions generated by Oracle XA applications, this column information applies specifically to the DBA_2PC_NEIGHBORS table:

- The DBID column is always xa_orcl
- The DBUSER_OWNER column is always *db_name*xa.oracle.com

Remember that the *db_name* is always specified as DB=*db_name* in the open string. If you do not specify this field in the open string, then the value of this column is NULLxa.oracle.com for transactions generated by Oracle XA applications.

For example, this SQL statement provide more information about in-doubt transactions generated by Oracle XA applications:

```
SELECT *
FROM DBA_2PC_PENDING p, DBA_2PC_NEIGHBORS n
WHERE p.LOCAL_TRAN_ID = n.LOCAL_TRAN_ID
AND n.DBID = 'xa_orcl';
```

Alternatively, if you know the format ID used by the transaction processing monitor, then you can use DBA_PENDING_TRANSACTIONS or V$GLOBAL_TRANSACTION. Whereas DBA_PENDING_TRANSACTIONS gives a list of prepared transactions, V$GLOBAL_TRANSACTION provides a list of all active global transactions.

Oracle XA Issues and Restrictions

Topics:

- Using Database Links in Oracle XA Applications
- Managing Transaction Branches in Oracle XA Applications
- Using Oracle XA with Oracle Real Application Clusters (Oracle RAC)
- SQL-Based Oracle XA Restrictions
- Miscellaneous Restrictions

Using Database Links in Oracle XA Applications

Oracle XA applications can access other Oracle Database instances through database links with these restrictions:

- They must use the shared server configuration.

  The transaction processing monitors (TPMs) use shared servers to open the connection to an Oracle Database A. Then the operating system network connection required for the database link is opened by the dispatcher instead of a dedicated server process. This allows different services or threads to operate on the transaction.

  If this restriction is not satisfied, then when you use database links within an XA transaction, it creates an operating system network connection between the dedicated server process and the other Oracle Database B. Because this network connection cannot be moved from one dedicated server process to another, you cannot detach from this dedicated server process of database A. Then when you access the database B through a database link, you receive an ORA-24777 error.

- The other database being accessed must be another Oracle Database.

If these restrictions are satisfied, Oracle Database allows such links and propagates the transaction protocol (prepare, rollback, and commit) to the other Oracle Database instances.

If using the shared server configuration is not possible, then access the remote database through the Pro*C/C++ application by using `EXEC SQL AT` syntax.

The `init.ora` parameter `OPEN_LINKS_PER_INSTANCE` specifies the number of open database link connections that can be migrated. These `dblink` connections are used by XA transactions so that the connections are cached after a transaction is committed. Another transaction can use the database link connection if the user who created the connection also created the transaction. This parameter is different from the `init.ora` parameter `OPEN_LINKS`, which specifies the maximum number of concurrent open connections (including database links) to remote databases in one session. The `OPEN_LINKS` parameter does not apply to XA applications.

Managing Transaction Branches in Oracle XA Applications

Oracle Database transaction branches within the same global transaction can be coupled tightly or loosely. If the transaction branches are **tightly coupled**, then they share locks. Consequently, pre-`COMMIT` updates in one transaction branch are visible in other branches that belong to the same global transaction. In loosely coupled transaction branches, the branches do not share locks and do not see updates in other branches.

In a tightly coupled branch, Oracle Database obtains the DX lock before running any statement. Because the system does not obtain a lock before running the statement, loosely coupled transaction branches result in greater concurrency. The disadvantage is that all transaction branches must go through the two phases of commit, that is, the system cannot use XA one-phase optimization.

Table 15–9 summarizes the trade-offs between tightly coupled branches and loosely coupled branches.

Table 15–9 Tightly and Loosely Coupled Transaction Branches

Attribute	Tightly Coupled Branches	Loosely Coupled Branches
Two Phase Commit	Read-only optimization	Two phases
	[prepare for all branches, commit for last branch]	[prepare and commit for all branches]
Serialization	Database call	None

Using Oracle XA with Oracle Real Application Clusters (Oracle RAC)

As of Oracle Database 11*g* Release 1 (11.1), an XA transaction can span Oracle RAC instances, allowing any application that uses XA to take full advantage of the Oracle RAC environment, enhancing the availability and scalability of the application.

> **Note:** External procedure callouts combined with distributed transactions is not supported.

Topics:

- GLOBAL_TXN_PROCESSES Initialization Parameter
- Managing Transaction Branches on Oracle RAC
- Managing Instance Recovery in Oracle RAC with DTP Services (10.2)

- Global Uniqueness of XIDs in Oracle RAC
- Tight and Loose Coupling

GLOBAL_TXN_PROCESSES Initialization Parameter

The initialization parameter `GLOBAL_TXN_PROCESSES` specifies the initial number of GTX*n* background processes for each Oracle RAC instance. Its default value is 1.

Leave this parameter at its default value clusterwide if distributed transactions might span multiple Oracle RAC instances. This allows the units of work performed across these Oracle RAC instances to share resources and act as a single transaction (that is, the units of work are tightly coupled). It also allows 2PC requests to be sent to any node in the cluster.

> **See Also:** *Oracle Database Reference* for more information about `GLOBAL_TXN_PROCESSES`

Managing Transaction Branches on Oracle RAC

> **Note:** This topic applies if either of the following is true:
>
> - The initialization parameter `GLOBAL_TXN_PROCESSES` is not at its default value in the initialization file of every Oracle RAC instance.
> - The Oracle XA application resumes or joins previously detached branches of a transaction.

Oracle Database permits different instances to operate on different transaction branches in Oracle RAC. For example, Node 1 can operate on branch A while Node 2 operates on branch B. Before Oracle Database 11*g* Release 1 (11.1), if transaction branches were on different instances, then they were loosely coupled and did not share locks. In this case, Oracle Database treated different units of work in different application threads as separate entities that did not share resources.

A different case is when multiple instances operate on a single transaction branch. For example, assume that a single transaction lands on Node 1 and Node 2 as follows:

Node 1

1. `xa_start`
2. SQL operations
3. `xa_end` (SUSPEND)

Node 2

1. `xa_start` (RESUME)
2. `xa_prepare`
3. `xa_commit`
4. `xa_end`

In the immediately preceding sequence, Oracle Database returns an error because Node 2 must not resume a branch that is physically located on a different node (Node 1).

Before Release 11.1, the way to achieve tight coupling in Oracle RAC was to use **Distributed Transaction Processing (DTP) services**, that is, services whose cardinality (one) ensured that all tightly-coupled branches landed on the same instance—regardless of whether load balancing was enabled. Middle-tier components addressed Oracle Database through a common logical database service name that mapped to a single Oracle RAC instance at any point in time. An intermediate name resolver for the database service hid the physical characteristics of the database instance. DTP services enabled all participants of a tightly-coupled global transaction to create branches on one instance.

As of Release 11.1, the DTP service is no longer required to support XA transactions with tightly coupled branches. By default, tightly coupled branches that land on different Oracle RAC instances remain tightly coupled; that is, they share locks and resources across Oracle RAC instances.

For example, when you use a DTP service, this sequence of actions occurs on the same instance:

1. `xa_start`
2. SQL operations
3. `xa_end` (SUSPEND)
4. `xa_start` (RESUME)
5. SQL operations
6. `xa_prepare`
7. `xa_commit` or `xa_rollback`

Moreover, multiple tightly-coupled branches land on the same instance if each addresses the Oracle RM with the same DTP service.

To leverage all instances in the cluster, create multiple DTP services, with one or more on each node that hosts distributed transactions. All branches of a global distributed transaction exist on the same instance. Thus, you can leverage all instances and nodes of an Oracle RAC cluster to balance the load of many distributed XA transactions, thereby maximizing application throughput.

> **See Also:** *Oracle Real Application Clusters Administration and Deployment Guide* to learn how to manage distributed transactions in a Real Application Clusters configuration

Managing Instance Recovery in Oracle RAC with DTP Services (10.2)

> **Note:** You need not read this topic if the initialization parameter `GLOBAL_TXN_PROCESSES` is at its default value in the initialization file of every Oracle RAC instance.

Before Oracle Database 10*g* Release 2 (10.2), TM was responsible for detecting failure and triggering failover and failback in Oracle RAC. To ensure that information about in-doubt transactions was propagated to `DBA_2PC_PENDING`, TM had to call `xa_recover` before resolving the in-doubt transactions. If an instance failed, then the XA client library could not fail over to another instance until it had run the`SYS.DBMS_XA.DIST_TXN_SYNC` procedure to ensure that the undo segments of the failed instance were recovered. As of Oracle Database 10*g* Release 2 (10.2), there is no such

requirement to call `xa_recover` in cases where the TM has enough information about in-flight transactions.

> **Note:** As of Oracle Database 9*g* Release 2 (9.2), `xa_recover` is required to wait for distributed data manipulation language (DML) statements to complete on remote sites.

Using DTP services in Oracle RAC has these benefits:

- Automates instance failure detection.
- Automates instance failover and failback. When an instance fails, the DTP service hosted on this instance fails over to another instance. The failover forces clients to reconnect; nevertheless, the logical names for the service remain the same. Failover is automatic and does not require an administrator intervention. The administrator can induce failback by a service relocate statement, but all failback-related recovery is automatically handled within the database server.
- Enables Oracle Database rather than the client to drive instance recovery. The database does not require middle-tier TM involvement to determine the state of transactions prepared by other instances.

> **See Also:** *Oracle Real Application Clusters Administration and Deployment Guide* to learn how to manage instance recovery

Global Uniqueness of XIDs in Oracle RAC

> **Note:** You need not read this topic if the initialization parameter `GLOBAL_TXN_PROCESSES` is at its default value in the initialization file of every Oracle RAC instance.

Before Oracle Database 11*g* Release 1 (11.1), Oracle RAC database cannot determine whether a given XID is unique for XA transactions throughout the cluster.

For example, suppose that there is an XID `Fmt(x).Tx(1).Br(1)` on Oracle RAC instance 1 and another XID `Fmt(x).Tx(1).Br(1)` on Oracle RAC instance 2. Each of these can start a branch and run SQL even though the XID is not unique across Oracle RAC instances.

As of Release 11.1, Oracle RAC database detects the duplicate XIDs across Oracle RAC instances and prevents a branch with a duplicate XID from starting.

Tight and Loose Coupling

> **Note:** You need not read this topic if the initialization parameter `GLOBAL_TXN_PROCESSES` is at its default value in the initialization file of every Oracle RAC instance.

Oracle Database transaction branches within the same global transaction can be coupled either tightly or loosely (for details, see "Managing Transaction Branches in Oracle XA Applications" on page 15-22). Ordinarily, coupling type is determined by the value of the `Loose_Coupling` field of the `xa_open` string (see Table 15–5 on

page 15-9). However, if branches are landed on different Oracle RAC instances when running Oracle RAC, they are loosely coupled even if `Loose_Coupling=false`.

SQL-Based Oracle XA Restrictions

This section describes restrictions concerning these SQL operations:

- Rollbacks and Commits
- DDL Statements
- Session State
- EXEC SQL

Rollbacks and Commits

Because the transaction manager is responsible for coordinating and monitoring the progress of the global transaction, the application must not contain any Oracle Database-specific statement that independently rolls back or commits a global transaction. However, you can use rollbacks and commits in a local transaction.

Do not use `EXEC SQL ROLLBACK WORK` for precompiler applications when you are in the middle of a global transaction. Similarly, an OCI application must not run `OCITransRollback`, or the Version 7 equivalent `orol`. You can roll back a global transaction by calling `tx_rollback`.

Similarly, a precompiler application must not have the `EXEC SQL COMMIT WORK` statement in the middle of a global transaction. An OCI application must not run `OCITransCommit` or the Version 7 equivalent `ocom`. For example, use `tx_commit` or `tx_rollback` to end a global transaction.

DDL Statements

Because a data definition language (DDL) statement, such as `CREATE TABLE`, implies an implicit commit, the Oracle XA application cannot run any DDL statements.

Session State

Oracle Database does not guarantee that session state is valid between TPM services. For example, if a TPM service updates a session variable (such as a global package variable), then another TPM service that runs as part of the same global transaction might not see the change. Use savepoints only within a TPM service. The application must not refer to a savepoint that was created in another TPM service. Similarly, an application must not attempt to fetch from a cursor that was executed in another TPM service.

EXEC SQL

Do not use the `EXEC SQL` statement to connect or disconnect. That is, do not use `EXEC SQL CONNECT`, `EXEC SQL COMMIT WORK RELEASE` or `EXEC SQL ROLLBACK WORK RELEASE`.

Miscellaneous Restrictions

- Oracle Database does not support association migration (a means whereby a transaction manager might resume a suspended branch association in another branch).
- The optional XA feature asynchronous XA calls is not supported.

- Set the `TRANSACTIONS` initialization parameter to the expected number of concurrent global transactions. The initialization parameter `OPEN_LINKS_PER_INSTANCE` specifies the number of open database link connections that can be migrated. These database link connections are used by XA transactions so that the connections are cached after a transaction is committed.

  See Also: "Using Database Links in Oracle XA Applications" on page 15-21

- The maximum number of `xa_open` calls for each thread is 32.
- When building an XA application based on TP-monitor, ensure that the TP-monitors libraries (that define the symbols `ax_reg` and `ax_unreg`) are placed in the link line before Oracle Database's client shared library. If your platform does not support shared libraries or if your linker is not sensitive to ordering of libraries in the link line, use Oracle Database's nonshared client library. These link restrictions are applicable only when using XA's dynamic registration (Oracle XA switch `xaoswd`).

16

Developing Applications with the Publish-Subscribe Model

This chapter explains how to develop applications on the publish-subscribe model.

Topics:

- Introduction to the Publish-Subscribe Model
- Publish-Subscribe Architecture
- Publish-Subscribe Concepts
- Examples of a Publish-Subscribe Mechanism

Introduction to the Publish-Subscribe Model

Because the database is the most significant resource of information within the enterprise, Oracle created a publish-subscribe solution for enterprise information delivery and messaging to complement this role.

Networking technologies and products enable a high degree of connectivity across a large number of computers, applications, and users. In these environments, it is important to provide asynchronous communications for the class of distributed systems that operate in a loosely-coupled and autonomous fashion, and which require operational immunity from network failures. This requirement is filled by various middleware products that are characterized as messaging, message-oriented middleware (MOM), message queuing, or publish-subscribe.

Applications that communicate through a publish and subscribe paradigm require the sending applications (publishers) to publish messages without explicitly specifying recipients or having knowledge of intended recipients. Similarly, receiving applications (subscribers) must receive only those messages that the subscriber has registered an interest in.

This decoupling between senders and recipients is usually accomplished by an intervening entity between the publisher and the subscriber, which serves as a level of indirection. This intervening entity is a queue that represents a subject or channel. Figure 16–1 illustrates publish and subscribe functionality.

Figure 16–1 Oracle Publish-Subscribe Functionality

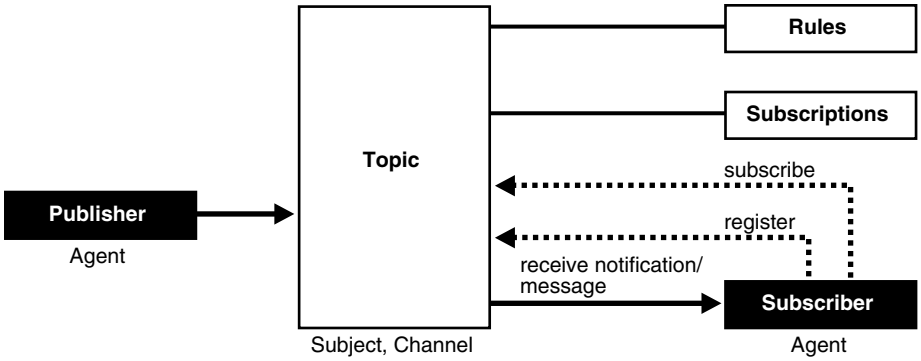

A subscriber subscribes to a queue by expressing interest in messages enqueued to that queue and by using a subject- or content-based rule as a filter. This results in a set of rule-based subscriptions associated with a given queue.

At run time, publishers post messages to various queues. The queue (in other words, the delivery mechanisms of the underlying infrastructure) then delivers messages that match the various subscriptions to the appropriate subscribers.

Publish-Subscribe Architecture

Oracle Database includes these features to support database-enabled publish-subscribe messaging:

- Database Events
- Oracle Advanced Queuing
- Client Notification

Database Events

Database events support declarative definitions for publishing database events, detection, and run-time publication of such events. This feature enables active publication of information to end-users in an event-driven manner, to complement the traditional pull-oriented approaches to accessing information.

See Also: *Oracle Database PL/SQL Language Reference*

Oracle Advanced Queuing

Oracle Advanced Queuing (AQ) supports a queue-based publish-subscribe paradigm. Database queues serve as a durable store for messages, along with capabilities to allow publish and subscribe based on queues. A rules-engine and subscription service dynamically route messages to recipients based on expressed interest. This allows decoupling of addressing between senders and receivers to complement the existing explicit sender-receiver message addressing.

See Also: *Oracle Streams Advanced Queuing User's Guide*

Client Notification

Client notifications support asynchronous delivery of messages to interested subscribers, enabling database clients to register interest in certain queues, and it enables these clients to receive notifications when publications on such queues occur.

Asynchronous delivery of messages to database clients is in contrast to the traditional polling techniques used to retrieve information.

See Also: *Oracle Call Interface Programmer's Guide*

Publish-Subscribe Concepts

queue

A **queue** is an entity that supports the notion of named subjects of interest. Queues can be characterized as persistent or nonpersistent (lightweight).

A **persistent queue** serves as a durable container for messages. Messages are delivered in a deferred and reliable mode.

The underlying infrastructure of a **nonpersistent, or lightweight, queue** pushes the messages published to connected clients in a lightweight, at-best-once, manner.

agent

Publishers and subscribers are internally represented as agents.

An **agent** is a persistent logical subscribing entity that expresses interest in a queue through a subscription. An agent has properties, such as an associated subscription, an address, and a delivery mode for messages. In this context, an agent is an electronic proxy for a publisher or subscriber.

client

A **client** is a transient physical entity. The attributes of a client include the physical process where the client programs run, the node name, and the client application logic. Several clients can act on behalf of a single agent. The same client, if authorized, can act on behalf of multiple agents.

rule on a queue

A **rule on a queue** is specified as a conditional expression using a predefined set of operators on the message format attributes or on the message header attributes. Each queue has an associated message content format that describes the structure of the messages represented by that queue. The message format may be unstructured (`RAW`) or it may have a well-defined structure (ADT). This allows both subject- or content-based subscriptions.

subscriber

Subscribers (agents) may specify subscriptions on a queue using a rule. Subscribers are durable and are stored in a catalog.

database event publication framework

The database represents a significant source for publishing information. An event framework is proposed to allow declarative definition of database event publication. As these pre-defined events occur, the framework detects and publishes such events. This allows active delivery of information to end-users in an event-driven manner as part of the publish-subscribe capability.

registration

Registration is the process of associated delivery information by a given client, acting on behalf of an agent. There is an important distinction between the subscription and registration related to the agent/client separation.

Subscription indicates an interest in a particular queue by an agent. It does not specify where and how delivery must occur. Delivery information is a physical property that is associated with a client, and it is a transient manifestation of the logical agent (the subscriber). A specific client process acting on behalf of an agent registers delivery information by associating a host and port, indicating *where* the delivery is to be done, and a callback, indicating *how* there delivery is to be done.

publishing a message

Publishers publish messages to queues by using the appropriate queuing interfaces. The interfaces may depend on which model the queue is implemented on. For example, an enqueue call represents the publishing of a message.

rules engine

When a message is posted or published to a given queue, a rules engine extracts the set of candidate rules from all rules defined on that queue that match the published message.

subscription services

Corresponding to the list of candidate rules on a given queue, the set of subscribers that match the candidate rules can be evaluated. In turn, the set of agents corresponding to this subscription list can be determined and notified.

posting

The queue notifies all registered clients of the appropriate published messages. This concept is called **posting**. When the queue must notify all interested clients, it posts the message to all registered clients.

receiving a message

A subscriber may receive messages through any of these mechanisms:

- A client process acting on behalf of the subscriber specifies a callback using the registration mechanism. The posting mechanism then asynchronously invokes the callback when a message matches the subscriber's subscription. The message content may be passed to the callback function (nonpersistent queues only).
- A client process acting on behalf of the subscriber specifies a callback using the registration mechanism. The posting mechanism then asynchronously invokes the callback function, but without the full message content. This serves as a notification to the client, which subsequently retrieves the message content in a pull fashion (persistent queues only).
- A client process acting on behalf of the subscriber simply retrieves messages from the queue in a periodic, or some other appropriate, manner. While the messages are deferred, there is no asynchronous delivery to the end-client.

Examples of a Publish-Subscribe Mechanism

This example shows how database events, client notification, and AQ work to implement publish-subscribe.

- Create under the user schema, `pubsub`, with all objects necessary to support a publish-subscribe mechanism. In this particular code, the Agent `snoop` subscribe to messages that are published at logon events. To use AQ functionality, user `pubsub` needs `AQ_ADMINISTRATOR_ROLE` privileges and `EXECUTE` privilege on `DBMS_AQ` and `DBMS_AQADM`.

```
Rem -----------------------------------------------------
REM create queue table for persistent multiple consumers:
Rem -----------------------------------------------------

Rem  Create or replace a queue table
BEGIN
DBMS_AQADM.CREATE_QUEUE_TABLE(
   Queue_table        =>  'Pubsub.Raw_msg_table',
   Multiple_consumers =>   TRUE,
   Queue_payload_type =>  'RAW',
   Compatible         =>  '8.1');
END;
/
Rem -----------------------------------------------------
Rem  Create a persistent queue for publishing messages:
Rem -----------------------------------------------------

Rem  Create a queue for logon events
BEGIN
   DBMS_AQADM.CREATE_QUEUE(
         Queue_name     =>   'Pubsub.Logon',
         Queue_table    =>   'Pubsub.Raw_msg_table',
         Comment        =>   'Q for error triggers');
END;
/

Rem -----------------------------------------------------
Rem  Start the queue:
Rem -----------------------------------------------------

BEGIN
   DBMS_AQADM.START_QUEUE('pubsub.logon');
END;
/

Rem -----------------------------------------------------
Rem  define new_enqueue for convenience:
Rem -----------------------------------------------------

CREATE OR REPLACE PROCEDURE New_enqueue(
               Queue_name      IN VARCHAR2,
               Payload         IN RAW ,
               Correlation     IN VARCHAR2 := NULL,
               Exception_queue IN VARCHAR2 := NULL)
AS

Enq_ct     DBMS_AQ.Enqueue_options_t;
Msg_prop   DBMS_AQ.Message_properties_t;
Enq_msgid  RAW(16);
Userdata   RAW(1000);

BEGIN
   Msg_prop.Exception_queue := Exception_queue;
   Msg_prop.Correlation := Correlation;
   Userdata := Payload;

DBMS_AQ.ENQUEUE(Queue_name, Enq_ct, Msg_prop, Userdata, Enq_msgid);
END;
/
```

```
Rem -----------------------------------------------------
Rem  add subscriber with rule based on current user name,
Rem  using correlation_id
Rem -----------------------------------------------------


DECLARE
Subscriber Sys.Aq$_agent;
BEGIN
   Subscriber := sys.aq$_agent('SNOOP', NULL, NULL);
DBMS_AQADM.ADD_SUBSCRIBER(
    Queue_name         => 'Pubsub.logon',
    Subscriber         => subscriber,
    Rule               => 'CORRID = ''HR'' ');
END;
/

Rem -----------------------------------------------------
Rem  create a trigger on logon on database:
Rem -----------------------------------------------------


Rem  create trigger on after logon:
CREATE OR REPLACE TRIGGER pubsub.Systrig2
   AFTER LOGON
   ON DATABASE
   BEGIN
      New_enqueue('Pubsub.Logon', HEXTORAW('9999'), Dbms_standard.login_user);
   END;
/
```

- After subscriptions are created, the next step is for the client to register for notification using callback functions. This is done using the Oracle Call Interface (OCI). This code performs necessary steps for registration. The initial steps of allocating and initializing session handles are omitted here for sake of clarity:

```
ub4 namespace = OCI_SUBSCR_NAMESPACE_AQ;

/* callback function for notification of logon of user 'HR' on database: */

ub4 notifySnoop(ctx, subscrhp, pay, payl, desc, mode)
dvoid *ctx;
OCISubscription *subscrhp;
dvoid *pay;
ub4 payl;
dvoid *desc;
ub4 mode;
{
    printf("Notification : User HR Logged on\n");
}

int main()
{
    OCISession *authp = (OCISession *) 0;
    OCISubscription *subscrhpSnoop = (OCISubscription *)0;

    /*****************************************************
       Initialize OCI Process/Environment
       Initialize Server Contexts
       Connect to Server
```

```
        Set Service Context
    ****************************************************/

    /* Registration Code Begins */

    /* Each call to initSubscriptionHn allocates
           and Initialises a Registration Handle */


    initSubscriptionHn(    &subscrhpSnoop,    /* subscription handle */
        "ADMIN:PUBSUB.SNOOP", /* subscription name */
                   /* <agent_name>:<queue_name> */
        (dvoid*)notifySnoop); /* callback function */

     /****************************************************
        The Client Process does not need a live Session for Callbacks
        End Session and Detach from Server
     ****************************************************/

    OCISessionEnd ( svchp,  errhp, authp, (ub4) OCI_DEFAULT);

    /* detach from server */
    OCIServerDetach( srvhp, errhp, OCI_DEFAULT);

    while (1)     /* wait for callback */
        sleep(1);

}

void initSubscriptionHn (subscrhp,
subscriptionName,
func)

OCISubscription **subscrhp;
char* subscriptionName;
dvoid * func;
{

    /* allocate subscription handle: */

    (void) OCIHandleAlloc((dvoid *) envhp, (dvoid **)subscrhp,
        (ub4) OCI_HTYPE_SUBSCRIPTION,
        (size_t) 0, (dvoid **) 0);

    /* set subscription name in handle: */

    (void) OCIAttrSet((dvoid *) *subscrhp, (ub4) OCI_HTYPE_SUBSCRIPTION,
        (dvoid *) subscriptionName,
        (ub4) strlen((char *)subscriptionName),
        (ub4) OCI_ATTR_SUBSCR_NAME, errhp);

    /* set callback function in handle: */

    (void) OCIAttrSet((dvoid *) *subscrhp, (ub4) OCI_HTYPE_SUBSCRIPTION,
        (dvoid *) func, (ub4) 0,
        (ub4) OCI_ATTR_SUBSCR_CALLBACK, errhp);

    (void) OCIAttrSet((dvoid *) *subscrhp, (ub4) OCI_HTYPE_SUBSCRIPTION,
        (dvoid *) 0, (ub4) 0,
        (ub4) OCI_ATTR_SUBSCR_CTX, errhp);
```

```
    /* set namespace in handle: */

    (void) OCIAttrSet((dvoid *) *subscrhp, (ub4) OCI_HTYPE_SUBSCRIPTION,
        (dvoid *) &namespace, (ub4) 0,
        (ub4) OCI_ATTR_SUBSCR_NAMESPACE, errhp);

    checkerr(errhp, OCISubscriptionRegister(svchp, subscrhp, 1, errhp,

        OCI_DEFAULT));
}
```

If user HR logs on to the database, the client is notified, and the call back function `notifySnoop` is invoked.

17

Using the Identity Code Package

The Identity Code Package is a feature in the Oracle Database that offers tools and techniques to store, retrieve, encode, decode, and translate between various product or identity codes, including Electronic Product Code (EPC), in an Oracle Database. The Identity Code Package provides data types, metadata tables and views, and PL/SQL packages for storing EPC standard RFID tags or new types of RFID tags in a user table.

The Identity Code Package empowers Oracle Database with the knowledge to recognize EPC coding schemes, support efficient storage and component level retrieval of EPC data, and comply with the EPCglobal Tag Data Translation 1.0 (TDT) standard that defines how to decode, encode, and translate between various EPC RFID tag representations.

The Identity Code Package also provides an extensible framework that allows developers to use pre-existing coding schemes with their applications that are not included in the EPC standard and make the Oracle Database adaptable to these older systems and to any evolving identity codes that may some day be part of a future EPC standard.

The Identity Code Package also lets developers create their own identity codes by first registering the encoding category, registering the encoding type, and then registering the components associated with each encoding type.

Topics.

- Identity Concepts
- What is the Identity Code Package?
- Using the Identity Code Package
- Identity Code Package Types
- DBMS_MGD_ID_UTL Package
- Identity Code Metadata Tables and Views
- Electronic Product Code (EPC) Concepts
- Oracle Database Tag Data Translation Schema

Identity Concepts

A database object `MGD_ID` is defined that lets users use EPC standard identity codes and use their own existing identity codes. See "Electronic Product Code (EPC) Concepts" on page 17-21 for a brief description of EPC concepts. The `MGD_ID` object serves as the base code object to which belong certain categories, or types of the RFID tag, such as the EPC category, NASA category, and many other categories. Each

category has a set of tag schemes or documents that define tag representation structures and their components. For the EPC category, the metadata needed to define encoding schemes (SGTIN-64, SGTIN-96, GID-96, and so forth) representing different encoding types (defined in the EPC standard v1.1) is loaded by default into the database. Users can define encoding their own categories and schemes as shown in Figure 17–1 and load these into the database as well.

Figure 17–1 RFID Code Categories and Their Schemes

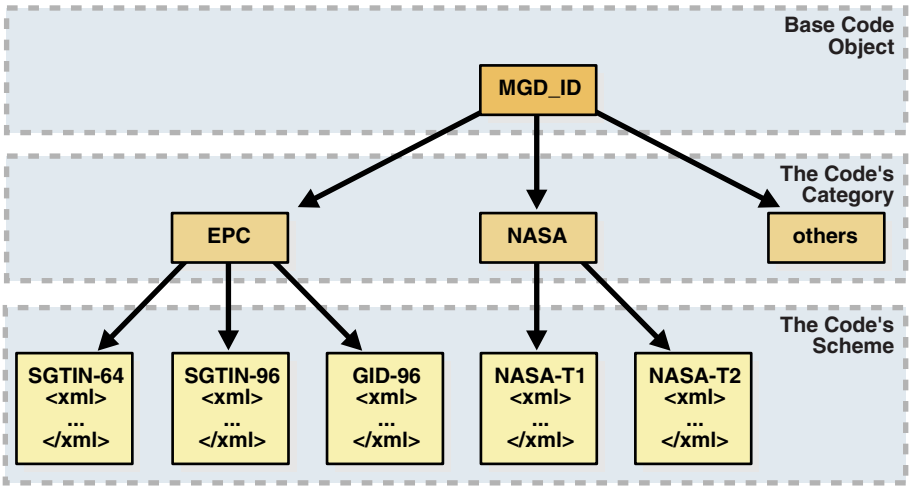

An `MGD_ID` object contains two attributes, a category_id and a list of components consisting of name-value pairs. When `MGD_ID` objects are stored, the tag representation must be parsed into these component name-value pairs upon object creation.

EPC standard version 1.1 defines one General Identifier type (GID) that is independent of any known, existing code schemes, five Domain Identifier types that are based on EAN.UCC specifications, and the identity type United States Department of Defense (USDOD). The five EAN.UCC based identity types are the serialized global trade identification number (SGTIN), the serial shipping container code (SSCC), the serialized global location number (SGLN), the global returnable asset identifier (GRAI) and the global individual asset identifier (GIAI).

Except GID, which has only one bit-level encoding, all the other identity types each have two encodings depending on their length: 64-bit and 96-bit. So in total there are thirteen different standard encodings for EPC tags. In addition, tags can be encoded in representations other than binary, such as the tag URI and pure identity representations.

Each EPC encoding has its own structure and organization, see Table 17–1. The EPC encoding structure field names relate to the names in the parameter_list parameter name-value pairs in the Identity Code Package API. For example, for SGTIN-64, the structure field names are Filter Value, Company Prefix Index, Item Reference, and Serial Number.

Table 17–1 General Structure of EPC Encodings

Encoding Name	Header Length in bits	Field Names (parameter_list name-value pairs) and (length in bits)
GID-96	8	General Manager Number (8), Object Class (24), Serial Number (36)

Table 17–1 (Cont.) General Structure of EPC Encodings

Encoding Name	Header Length in bits	Field Names (parameter_list name-value pairs) and (length in bits)
SGTIN-64	2	Filter Value (3), Company Prefix Index (14), Item Reference 20), Serial Number (25)
SGTIN-96	8	Filter Value (3), Partition (3), Company Prefix (20-40), Item Reference (24-4), Serial Number (38)
SSCC-64	8	Filter Value (3), Company Prefix Index (14), Serial Reference (39)
SSCC-96	8	Filter Value (3), Partition (3), Company Prefix (20-40), Serial Reference (38-18), Unallocated (24)
SGLN-64	8	Filter Value (3), Company Prefix Index (14), Location Reference (20), Serial Number (19)
SGLN-96	8	Filter Value (3), Partition (3), Company Prefix (20-40), Location Reference (21-1), Serial Number (41)
GRAI-64	8	Filter Value (3), Company Prefix Index (14), Asset Type (20), Serial Number (19)
GRAI-96	8	Filter Value (3), Partition (3), Company Prefix (20-40), Asset Type (24-4), Serial Number (38)
GIAI-64	8	Filter Value (3), Company Prefix Index (14), Individual Asset Reference (39)
GIAI-96	8	Filter Value (3), Partition (3), Company Prefix (20-40), Individual Asset Reference (62-42)
USDOD-64	8	Filter Value (2), Government Managed Identifier (30), Serial Number (24)
USDOD-96	8	Filter Value (4), Government Managed Identifier (48), Serial Number (36)

EPCglobal defines eleven tag schemes (GID-96, SGTIN-64, SGTIN-96, and so forth). Each of these schemes has various representations; today, the most often used are BINARY, TAG_URI, and PURE_IDENTITY. For example, information in an SGTIN-64 can be represented in these ways:

```
BINARY: 1001100000000000001000001110110001000010000011111110011000110010
PURE_IDENTITY:  urn:epc:id:sgtin:0037000.030241.1041970
TAG_URI: urn:epc:tag:sgtin-64:3.0037000.030241.1041970
LEGACY: gtin=00037000302414;serial=1041970
ONS_HOSTNAME: 030241.0037000.sgtin.id.example.com
```

Some representations contain all information about the tag (BINARY and TAG_URI), while other representations contain only partial information (PURE_IDENTITY). It is therefore possible to translate a tag from its TAG_URI to its PURE_IDENTITY representation, but it is not possible to translate in the other direction without more information being provided, namely the filter value must be supplied.

EPCglobal released a Tag Data Translation 1.0 (TDT) standard that defines how to decode, encode, and translate between various EPC RFID tag representations. Decoding refers to parsing a given representation into field/value pairs, and encoding refers to reconstructing representations from these fields. Translating refers to decoding one representation and instantly encoding it into another.TDT defines this information using a set of XML files, each referred to as a scheme. For example, the SGTIN-64 scheme defines how to decode, encode, and translate between various

SGTIN-64 representations, such as binary and pure identity. For details about the EPCglobal TDT schema, see the EPCglobal Tag Data Translation specification.

A key feature of the TDT specification is its ability to define any EPC scheme using the same XML schema. This approach creates a standard way of defining EPC metadata that RFID applications can then use to write their parsers, encoders, and translators. When the application is written according to the TDT specification, it must be able to update its set of EPC tag schemes and modify its action according to the metadata.

The Oracle Database metadata structure is similar, but not identical to the TDT standard. To fit the EPCglobal TDT specification, the Oracle RFID package must be able to ingest any TDT compatible scheme and seamlessly translate it into the generic Oracle Database defined metadata. See the `EPC_TO_ORACLE` Function in Table 17–4 for more information.

Reconstructing tag representation from fields, or in other words, encoding tag data into predefined representations is easily accomplished using the `MGD_ID.format` function. Likewise, the decoding of tag representations into `MGD_ID` objects and then encoding these objects into tag representations is also easily accomplished using the `MGDID.translate` function. See the `FORMAT` Member Function and the `TRANSLATE` Static Function in Table 17–3 for more information.

Because the EPCglobal TDT standard is powerful and highly extensible, the Oracle RFID standard metadata is a close relative of the TDT specification. See "Oracle Database Tag Data Translation Schema" on page 17-24 for the actual Oracle Database TDT XML schema. Developers can refer to this Oracle Database TDT XML schema to define their own tag structures.

Figure 17–2 shows the Oracle Database Tag Data Translation Markup Language Schema diagram.

Figure 17–2 Oracle Database Tag Data Translation Markup Language Schema

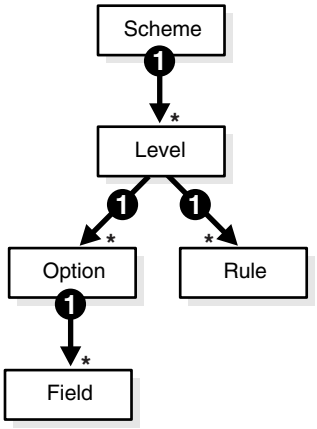

The top level element in a tag data translation xml is 'scheme'. Each scheme defines various tag encoding representations, or levels. SGTIN-64 and GID-96 are examples of tag encoding schemes, and BINARY or PURE_IDENTITY are examples of levels within these schemes. Each level has a set of options that define how to parse various representations into fields, and rules that define how to derive values for fields that require additional work, such as an external table lookup or the concatenation of other parsed out fields. See the EPCGlobal Tag Translator Specification for more information.

What is the Identity Code Package?

The Identity Code Package provides an extensible framework that supports the current RFID tags with the standard family of EPC bit encodings for the supported encoding types and new and evolving tag encodings that are not included in the current EPC standard.

The Identity Code Package defines these ADTs:

- `MGD_ID` -- defines these (see `MGD_ID` ADT in Table 17–2 for more information):
  - Two attributes, `category_id` and `components`.
  - Four `MGD_ID` constructor functions for constructing identity code type objects to represent RFID tags.
  - A set of member subprograms for operating on these ADTs.

  "Using the Identity Code Package" on page 17-6 describes how to use these ADTs and member functions.

  "Identity Code Package Types" on page 17-18 and "DBMS_MGD_ID_UTL Package" on page 17-18 briefly describe the reference information for these ADTs along with a set of utility subprograms. See *Oracle Database PL/SQL Packages and Types Reference* for detailed reference information.
- `MGD_ID_COMPONENT` — defines two attributes, `comp_name`, which identifies the name of the component and `comp_value`, which identifies the components value.
- `MGD_ID_COMPONENT_VARRAY` — defines an array type that can store up to 128 elements of `MGD_IDCOMPONENT` type, which is used in two constructor functions for creating an identity code type object with a list of components.

The Identity Code Package supports EPC spec v1.1 by supplying the predefined `EPC_ENCODING_CATEGORY` encoding_category attribute definition with its bit-encoding structures for the supported encoding types. This information is stored as meta information in the supplied encoding metadata views, `MGD_USR_ID_CATEGORY`, `MGD_USR_ID_SCHEME`, the read-only views `MGD_ID_CATEGORY`, `MGD_ID_SCHEME`, and their underlying tables: `MGD_ID_CATEGORY_TAB`, `MGD_ID_SCHEME_TAB`, `MGD_ID_XML_VALIDATOR`. See these topics and files for more information:

- "Electronic Product Code (EPC) Concepts" on page 17-21 describes the EPC spec v1.1 product code and its family of coding schemes.
- "Identity Code Metadata Tables and Views" on page 17-19 describes the structure of the identity code meta tables and views and how metadata are used by the Identity Code Package to interpret the various RFID tags.
- The `mgdmeta.sql` file describes the meta table data for the `EPC_ENCODING_CATEGORY` categories and each of its specific encoding schemes.

After storing many thousands of RFID tags into the column of `MGD_ID` column type of your user table, you can improve query performance by creating an index on this column. See these topics for more information:

- "Building a Function-Based Index Using the Member Functions of the MGD_ID Column Type" on page 17-10 describes how to create a function based index or bitmap function based index using the member functions of the `MGD_ID` ADT.

The Identity Code Package provides a utility package that consists of various utility subprograms. See this topic for more information:

- "Identity Code Package Types" on page 17-18 and "DBMS_MGD_ID_UTL Package" on page 17-18 describes each of the member subprograms. A proxy

utility seta and removea proxy information. A metadata utility geta a category ID, refreshes a tag scheme for a category, removes a tag scheme for a category, and validates a tag scheme. A conversion utility translates standard EPCglobal Tag Data Translation (TDT) files into Oracle Database TDT files.

The Identity Code Package is extensible and lets you create your own identity code types for your new or evolving RFID tags. You can define your identity code types, `catagory_id` attribute values, and components structures for your own encoding types. See these topics for more information:

- "Creating a Category of Identity Codes" on page 17-13 describes how to create your own identity codes by first registering the encoding category, and then registering the schemes associated to the encoding category.
- "Identity Code Metadata Tables and Views" on page 17-19 describes the structure of the identity code meta tables and views and how to register meta information by storing it in the supplied metadata tables and views.

Using the Identity Code Package

Topics:

- Storing RFID Tags in Oracle Database Using MGD_ID ADT
- Building a Function-Based Index Using the Member Functions of the MGD_ID Column Type
- Using MGD_ID ADT Functions
- Defining a Category of Identity Codes and Adding Encoding Schemes to an Existing Category

Storing RFID Tags in Oracle Database Using MGD_ID ADT

Topics:

- Creating a Table with MGD_ID Column Type and Storing EPC Tag Encodings in the Column
- Constructing MGD_ID Objects to Represent RFID Tags
- Inserting an MGD_ID Object into a Database Table
- Querying MGD_ID Column Type

Creating a Table with MGD_ID Column Type and Storing EPC Tag Encodings in the Column

You can create tables using `MGD_ID` as the column type to represent RFID tags, for example:

Example 1. Using the `MGD_ID` column type:

```
CREATE TABLE Warehouse_info (
             Code          MGD_ID,
             Arrival_time  TIMESTAMP,
             Location      VARCHAR2(256);
             ...);
```

SQL*Plus command:

```
describe warehouse_info;
```

Result:

```
Name                                      Null?    Type
----------------------------------------- -------- ----------------------------
CODE                                      NOT NULL MGDSYS.MGD_ID
ARRIVAL_TIME                                       TIMESTAMP(6)
LOCATION                                           VARCHAR2(256)
```

Constructing MGD_ID Objects to Represent RFID Tags

There are several ways to construct `MGD_ID` objects:

- Constructing an MGD_ID Object (SGTIN-64) Passing in the Category ID and a List of Components
- Constructing an MGD_ID object (SGTIN-64) and Passing in the Category ID, the Tag Identifier, and the List of Additional Required Parameters
- Constructing an MGD_ID object (SGTIN-64) and Passing in the Category Name, Category Version (if null, then the latest version is used), and a List of Components
- Constructing an MGD_ID object (SGTIN-64) and Passing in the Category Name and Category Version, the Tag Identifier, and the List of Additional Required Parameters

Constructing an MGD_ID Object (SGTIN-64) Passing in the Category ID and a List of Components If a RFID tag complies to the EPC standard, an `MGD_ID` object can be created using its category ID and a list of components. For example:

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');
call DBMS_MGD_ID_UTL.refresh_category('1');
select MGD_ID ('1',
               MGD_ID_COMPONENT_VARRAY(
               MGD_ID_COMPONENT('companyprefix','0037000'),
               MGD_ID_COMPONENT('itemref','030241'),
               MGD_ID_COMPONENT('serial','1041970'),
               MGD_ID_COMPONENT('schemes','SGTIN-64')
              )
             ) from DUAL;
call DBMS_MGD_ID_UTL.remove_proxy();

@constructor11.sql
.
.
.
MGD_ID ('1', MGD_ID_COMPONENT_VARRAY
        (MGD_ID_COMPONENT('companyprefix', '0037000'),
        MGD_ID_COMPONENT('itemref', '030241'),
        MGD_ID_COMPONENT('serial', '1041970'),
        MGD_ID_COMPONENT('schemes', 'SGTIN-64')))
.
.
.
```

Constructing an MGD_ID object (SGTIN-64) and Passing in the Category ID, the Tag Identifier, and the List of Additional Required Parameters Use this constructor when there is a list of additional parameters required to create the `MGD_ID` object. For example:

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');
call DBMS_MGD_ID_UTL.refresh_category('1');
```

```
select MGD_ID('1',
              'urn:epc:id:sgtin:0037000.030241.1041970',
              'filter=3;scheme=SGTIN-64') from DUAL;
call DBMS_MGD_ID_UTL.remove_proxy();


@constructor22.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('filter', '3'),
       MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
       MGD_ID_COMPONENT('companyprefixlength', '7'),
       MGD_ID_COMPONENT('companyprefix', '0037000'),
       MGD_ID_COMPONENT('scheme', 'SGTIN-64'),
       MGD_ID_COMPONENT('serial', '1041970'),
       MGD_ID_COMPONENT('itemref', '030241')))
.
.
.
```

Constructing an MGD_ID object (SGTIN-64) and Passing in the Category Name, Category Version (if null, then the latest version is used), and a List of Components Use this constructor when a category version must be specified along with a category ID and a list of components. For example:

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');
call DBMS_MGD_ID_UTL.refresh_category
  (DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
select MGD_ID('EPC', NULL,
              MGD_ID_COMPONENT_VARRAY(
              MGD_ID_COMPONENT('companyprefix','0037000'),
              MGD_ID_COMPONENT('itemref','030241'),
              MGD_ID_COMPONENT('serial','1041970'),
              MGD_ID_COMPONENT('schemes','SGTIN-64')
             )
            ) from DUAL;
call DBMS_MGD_ID_UTL.remove_proxy();

@constructor33.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY
             (MGD_ID_COMPONENT('companyprefix', '0037000'),
              MGD_ID_COMPONENT('itemref', '030241'),
              MGD_ID_COMPONENT('serial', '1041970'),
              MGD_ID_COMPONENT('schemes', 'SGTIN-64')
             )
       )
.
.
.
```

Constructing an MGD_ID object (SGTIN-64) and Passing in the Category Name and Category Version, the Tag Identifier, and the List of Additional Required Parameters Use this constructor when the category version and an additional list of parameters is required.

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');
call DBMS_MGD_ID_UTL.refresh_category
```

```
  (DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
select MGD_ID('EPC', NULL,
              'urn:epc:id:sgtin:0037000.030241.1041970',
              'filter=3;scheme=SGTIN-64') from DUAL;
call DBMS_MGD_ID_UTL.remove_proxy();

@constructor44.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY
       (MGD_ID_COMPONENT('filter', '3'),
        MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
        MGD_ID_COMPONENT('companyprefixlength', '7'),
        MGD_ID_COMPONENT('companyprefix', '0037000'),
        MGD_ID_COMPONENT('scheme', 'SGTIN-64'),
        MGD_ID_COMPONENT('serial', '1041970'),
        MGD_ID_COMPONENT('itemref', '030241')
       )
      )
.
.
.
```

Inserting an MGD_ID Object into a Database Table

This example shows how to populate the `WAREHOUSE_INFO` table by inserting each `MGD_ID` object into the table along with the additional column values:

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');

call DBMS_MGD_ID_UTL.refresh_category
  (DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));

INSERT INTO WAREHOUSE_INFO (code, arrival_time, location)
   values (MGDSYS.MGD_ID ('EPC',
                          NULL,
                          'urn:epc:id:sgtin:0037000.030241.1041970',
                          null
                         ),
           SYSDATE,
           'SHELF_123');

INSERT INTO WAREHOUSE_INFO (code, arrival_time, location)
  values (MGDSYS.MGD_ID ('EPC',
                         NULL,
                         'urn:epc:id:sgtin:0037000.053021.1012353',
                         null
                        ),
          SYSDATE,
          'SHELF_456');
INSERT INTO WAREHOUSE_INFO (code, arrival_time, location)
  values (MGDSYS.MGD_ID ('EPC',
                         NULL,
                         'urn:epc:id:sgtin:0037000.020140.10174832',
                         null
                        ),
          SYSDATE,
          'SHELF_1034');

COMMITT;
```

```
call DBMS_MGD_ID_UTL.remove_proxy();
```

Querying MGD_ID Column Type

There are three ways to query on MGD_ID column type.

- Query the MGD_ID column type. Find all items with item reference 030241.

```
SELECT location, wi.code.get_component('itemref') as itemref,
                 wi.code.get_component('serial') as serial
FROM warehouse_info wi WHERE wi.code.get_component('itemref') = '030241';

LOCATION       |ITEMREF   |SERIAL
---------------|----------|----------
SHELF_123      |030241    |1041970
```

- Query using the member functions of the MGD_ID ADT. Select the pure identity representations of all RFID tags in the table.

```
SELECT wi.code.format(null,'PURE_IDENTITY')
   as PURE_IDENTITY FROM warehouse_info wi;

PURE_IDENTITY
--------------------------------------------------------------------------------
urn:epc:id:sgtin:0037000.030241.1041970
urn:epc:id:gid:0037000.053021.1012353
urn:epc:id:sgtin:0037000.020140.10174832
```

  See "Using the get_component Function with the MGD_ID Object" on page 17-11 for more information and see Table 17–3 for a list of member functions.

Building a Function-Based Index Using the Member Functions of the MGD_ID Column Type

You can improve the performance of queries based on a certain component of the RFID tags by creating a function-based index that uses the get_component member function or its variation convenience functions. For example:

```
CREATE INDEX warehouseinfo_idx2
  on warehouse_info(code.get_component('itemref'));
```

You can also improve the performance of queries based on a certain component of the RFID tags by creating a bitmap function based index that uses the get_component member function or its variation convenience functions. For example:

```
CREATE BITMAP INDEX warehouseinfo_idx3
  on warehouse_info(code.get_component('serial'));
```

Using MGD_ID ADT Functions

The MGD_ID ADT contains member subprograms that operate on these ADTs. See Table 17–2 for MGD_ID_COMPONENT, MGD_ID_COMPONENT_VARRAY, MGD_ID ADT reference information. See the mgdtyp.sql file for the MGD_ID ADT definition and its member subprograms.

Topics:

- Using the get_component Function with the MGD_ID Object
- Parsing Tag Data from Standard Representations
- Reconstructing Tag Representations from Fields

- Translating Between Tag Representations

Using the get_component Function with the MGD_ID Object

The `get_component` function is defined as follows:

```
MEMBER FUNCTION get_component(component_name IN VARCHAR2)
   RETURN VARCHAR2 DETERMINISTIC,
```

Each component in a identity code has a name. It is defined when the code type is registered. See "Defining a Category of Identity Codes and Adding Encoding Schemes to an Existing Category" on page 17-13 for more information about how to create a identity code type.

The `get_component` function takes the name of the component, `component_name` as a parameter, uses the metadata registered in the metadata table to analyze the identity code, and returns the component with the name `component_name`.

The `get_component` function can be used in a SQL query. For example, find the current location of the coded item for the component named `itemref`; or, in other words find all items with the item reference of 03024. Because the code tag has encoded the "itemref" as a component, you can use this SQL query:

```
SELECT location,
       w.code.get_component('itemref') as itemref,
       w.code.get_component('serial')  as serial
FROM   warehouse_info w
       WHERE  w.code.get_component('itemref')  = '030241';

LOCATION       |ITEMREF   |SERIAL
---------------|----------|----------
SHELF_123      |030241    |1041970
```

See Table 17–3 for a list of other member functions.

Parsing Tag Data from Standard Representations

RFID readers read the bit strings stored in the tags. The tag data and other information, such as the reader ID and the time stamp, first go through an edge server to be processed, normalized, and preliminarily filtered. Then, in many application scenarios, the information must be persistently stored and later on be retrieved. The Oracle Database understands the code structures representations of various EPC tags as described in Table 17–1 because these code representation schemes defined in the EPC Standard are pre-registered. This gives the Oracle Database the ability to understand all the EPC code schemes and parse various tag representations into fields. Users can also register their own coding structures for the identity codes that use other encoding technologies. In this way the system is extensible.

As mentioned in "Identity Concepts" on page 17-1, each of the EPCGlobal tag schemes (GID-96, SGTIN-64, SGTIN-96, and so forth) has various representations with the most often used ones being BINARY, TAG_URI, and PURE_IDENTITY.

Some representations contain all the information about the tag (BINARY and TAG_URI), while representations contain only partial information (PURE_IDENTITY). It is therefore possible to translate a tag from it's TAG_URI to it's PURE_IDENTITY representation, but it is not possible to translate in the other direction (PURE_IDENTITY to TAG_URI) without supplying more information, namely the filter value.

One MGD_ID constructor takes in four fields, the category name (such as EPC), the category version, the tag identifier (for EPC, the identifier must be in a representation previously described), and a parameter list for any additional parameters that may be

required to parse the tag representation. For example, this code creates an MGD_ID object from its BINARY representation.

```
SELECT MGD_ID
   ('EPC',
    null,
    '1001100000000000001000001110110001000010000011111110011000110010',
    null
   )
   AS NEW_RFID_CODE FROM DUAL;

NEW_RFID_CODE(CATEGORY_ID, COMPONENTS(NAME, VALUE))
--------------------------------------------------------------------------------
MGD_ID ('1',
        MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('filter', '3'),
        MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
        MGD_ID_COMPONENT('companyprefixlength', '7'),
        MGD_ID_COMPONENT('companyprefix', '0037000'),
        MGD_ID_COMPONENT('companyprefixindex', '1'),
        MGD_ID_COMPONENT('serial', '1041970'),
        MGD_ID_COMPONENT('itemref', '030241')
       )
      )
```

For example, an identical object can be created if the call is done with the TAG_URI representation of the tag as follows with the addition of the value of the filter value:

```
SELECT MGD_ID ('EPC',
               null,
               'urn:epc:tag:sgtin-64:3.0037000.030241.1041970',
               null
              )
  as NEW_RFID_CODE FROM DUAL;

NEW_RFID_CODE(CATEGORY_ID, COMPONENTS(NAME, VALUE))
--------------------------------------------------------------------------------
MGD_ID ('1',
        MGD_ID_COMPONENT_VARRAY (
         ( MGD_ID_COMPONENT('filter', '3'),
           MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
           MGD_ID_COMPONENT('companyprefixlength', '7'),
           MGD_ID_COMPONENT('companyprefix', '0037000'),
           MGD_ID_COMPONENT('serial', '1041970'),
           MGD_ID_COMPONENT('itemref', '030241')
         )
       )
```

Reconstructing Tag Representations from Fields

Another useful feature of the Identity Code package is the ability to encode tag data into predefined representations. For example, a warehouse wants to send certain inventory to a retailer, but first it wants to send an invoice that tells the retailer what inventory to expect. The invoice can be a list of pure identity URIs that the warehouse intends to send. If all the inventory in the `WAREHOUSE_INFO` table is to be sent, this example constructs the desired URIs:

```
SELECT wi.code.format (null,'PURE_IDENTITY')
  as PURE_IDENTITY FROM warehouse_info wi;

PURE_IDENTITY
--------------------------------------------------------------------------------
```

```
urn:epc:id:sgtin:0037000.030241.1041970
urn:epc:id:gid:0037000.053021.1012353
urn:epc:id:sgtin:0037000.020140.10174832
```

Translating Between Tag Representations

The Identity Code package can decode tag representations into MGD_ID objects and encode these objects into tag representations. These two steps can be combined into one step using the `MGD_ID.translate` function. Static translation allows for the conversion of an RFID tag from one representation to another. For example:

```
SELECT MGD_ID.translate ('EPC',
                         null,
                         'urn:epc:id:sgtin:0037000.030241.1041970',
                         'filter=3;scheme=SGTIN-64',
                         'BINARY'
                        )
  as BINARY FROM DUAL;

BINARY
--------------------------------------------------------------------------------
10011000000000000010000011101100010000100000111111100110001100 10
```

In this example, the binary representation contains more information than the pure identity representation. Specifically, it also contains the filter value and in this case the scheme value must also be specified to distinguish SGTIN-64 from SGTIN-96. Thus, the function call must provide the missing filter parameter information and specify the scheme name in order for translation call to succeed.

Defining a Category of Identity Codes and Adding Encoding Schemes to an Existing Category

Topics:

- Creating a Category of Identity Codes
- Adding Two Metadata Schemes to a Newly Created Category

Creating a Category of Identity Codes

Because the EPCglobal TDT standard is powerful and highly extensible, the Oracle Database RFID standard metadata is a close relative of the TDT specification. Thus, the Identity Code package is extensible: You can create your own categories and tag structures using generic metadata. To create a category of identity codes, use the `DBMS_MGD_ID_UTIL.create_category` function.

For example, suppose you want to create a category called MGD_SAMPLE_CATEGORY, which has two types of tags, a CONTRACTOR_TAG and an EMPLOYEE_TAG. This category and its two metadata schemes might be used within a company that must grant different access privileges to people who are full time employees from those who are contractors, and thus require that their security software be able to identify quickly between the two badge types at an RFID reader. This script creates a category named 'MGD_SAMPLE_CATEGORY', with a 1.0 category version, having an agency name as Oracle, with a URI as `http://www.oracle.com/mgd/sample`. See "Adding Two Metadata Schemes to a Newly Created Category" on page 17-14 for an example.

Adding Two Metadata Schemes to a Newly Created Category

Next, create an CONTRACTOR_TAG metadata scheme such as:

```
<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="CONTRACTOR_TAG" optionKey="1" xmlns="">
  <level type="URI" prefixMatch="mycompany.contractor.">
   <option optionKey="1" pattern="mycompany.contractor.([0-9]*).([0-9]*)"
           grammar="''mycompany.contractor.'' contractorID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="contractorID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="11">
   <option optionKey="1" pattern="11([01]{7})([01]{6})"
           grammar="''11'' contractorID divisionID ">
    <field seq="1" characterSet="[01]*" name="contractorID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
</TagDataTranslation>
```

The CONTRACTOR_TAG scheme contains two encoding levels, or ways in which the tag can be represented. The first level is URI and the second level is BINARY. The URI representation starts with the prefix "mycompany.contractor." and is then followed by two numeric fields separated by a period. The names of the two fields are contractorID and divisionID. The pattern field in the option tag defines the parsing structure of the tag URI representation, and the grammar field defines how to reconstruct the URI representation. The BINARY representation can be understood in a similar fashion. This representation starts with the prefix "01" and is then followed by the same two fields, contractorID and divisionID, this time, in their respective binary formats. Given this XML metadata structure, contractor tags can now be decoded from their URI and BINARY representations and the resulting fields can be re-encoded into one of these representations.

The EMPLOYEE_TAG scheme is defined in a similar fashion and is shown as follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="EMPLOYEE_TAG" optionKey="1" xmlns="">
  <level type="URI" prefixMatch="mycompany.employee.">
   <option optionKey="1" pattern="mycompany.employee.([0-9]*).([0-9]*)"
           grammar="''mycompany.employee.'' employeeID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="employeeID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="01">
   <option optionKey="1" pattern="01([01]{7})([01]{6})"
           grammar="''01'' employeeID divisionID ">
    <field seq="1" characterSet="[01]*" name="employeeID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
```

```
</TagDataTranslation>;
```

To add these schemes to the category ID previously created, use the DBMS_MGD_ID_UTIL.add_scheme function.

This script creates the MGD_SAMPLE_CATEGORY category, adds a contractor scheme and an employee scheme to the MGD_SAMPLE_CATEGORY category, validates the MGD_SAMPLE_CATEGORY scheme, tests the tag translation of the contractor scheme and the employee scheme, then removes the contractor scheme, tests the tag translation of the contractor scheme and this returns the expected exception for the removed contractor scheme, tests the tag translation of the employee scheme and this returns the expected values, then removes the MGD_SAMPLE_CATEGORY category:

```
--contents of add_scheme2.sql
SET LINESIZE 160
CALL DBMS_MGD_ID_UTL.set_proxy('www-proxy.us.oracle.com', '80');
-----------------------------------------------------------------------
---CREATE CATEGORY, ADD_SCHEME, REMOVE_SCHEME, REMOVE_CATEGORY-------
-----------------------------------------------------------------------
DECLARE
  amt          NUMBER;
  buf          VARCHAR2(32767);
  pos          NUMBER;
  tdt_xml      CLOB;
  validate_tdtxml VARCHAR2(1042);
  category_id  VARCHAR2(256);
BEGIN
  -- remove the testing category if it exists
  DBMS_MGD_ID_UTL.remove_category('MGD_SAMPLE_CATEGORY', '1.0');
  -- create the testing category 'MGD_SAMPLE_CATEGORY', version 1.0
  category_id := DBMS_MGD_ID_UTL.CREATE_CATEGORY('MGD_SAMPLE_CATEGORY', '1.0', 'Oracle',
'http://www.oracle.com/mgd/sample');
  -- add contractor scheme to the category
  DBMS_LOB.CREATETEMPORARY(tdt_xml, true);
  DBMS_LOB.OPEN(tdt_xml, DBMS_LOB.LOB_READWRITE);

  buf := '<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="CONTRACTOR_TAG" optionKey="1" xmlns="">
  <level type="URI" prefixMatch="mycompany.contractor.">
   <option optionKey="1" pattern="mycompany.contractor.([0-9]*).([0-9]*)"
           grammar="''mycompany.contractor.'' contractorID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="contractorID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="11">
   <option optionKey="1" pattern="11([01]{7})([01]{6})"
           grammar="''11'' contractorID divisionID ">
    <field seq="1" characterSet="[01]*" name="contractorID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
</TagDataTranslation>';

  amt := length(buf);
  pos := 1;
```

```
  DBMS_LOB.WRITE(tdt_xml, amt, pos, buf);
  DBMS_LOB.CLOSE(tdt_xml);

  DBMS_MGD_ID_UTL.ADD_SCHEME(category_id, tdt_xml);

  -- add employee scheme to the category
  DBMS_LOB.CREATETEMPORARY(tdt_xml, true);
  DBMS_LOB.OPEN(tdt_xml, DBMS_LOB.LOB_READWRITE);

  buf := '<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="EMPLOYEE_TAG" optionKey="1" xmlns="">
  <level type="URI" prefixMatch="mycompany.employee.">
   <option optionKey="1" pattern="mycompany.employee.([0-9]*).([0-9]*)"
           grammar="''mycompany.employee.'' employeeID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="employeeID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="01">
   <option optionKey="1" pattern="01([01]{7})([01]{6})"
           grammar="''01'' employeeID divisionID ">
    <field seq="1" characterSet="[01]*" name="employeeID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
</TagDataTranslation>';

  amt := length(buf);
  pos := 1;
  DBMS_LOB.WRITE(tdt_xml, amt, pos, buf);
  DBMS_LOB.CLOSE(tdt_xml);
  DBMS_MGD_ID_UTL.ADD_SCHEME(category_id, tdt_xml);

  -- validate the scheme
  dbms_output.put_line('Validate the MGD_SAMPLE_CATEGORY Scheme');
  validate_tdtxml := DBMS_MGD_ID_UTL.validate_scheme(tdt_xml);
  dbms_output.put_line(validate_tdtxml);
  dbms_output.put_line('Length of scheme xml is: '||DBMS_LOB.GETLENGTH(tdt_xml));

  -- test tag translation of contractor scheme
  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     'mycompany.contractor.123.45',
                     NULL, 'BINARY'));

  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     '111111011101101',
                     NULL, 'URI'));

  -- test tag translation of employee scheme
  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     'mycompany.employee.123.45',
                     NULL, 'BINARY'));
```

```
  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     '011111011101101',
                     NULL, 'URI'));

  DBMS_MGD_ID_UTL.REMOVE_SCHEME(category_id, 'CONTRACTOR_TAG');

  -- Test tag translation of contractor scheme. Doesn't work any more.
  BEGIN
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                       'mycompany.contractor.123.45',
                       NULL, 'BINARY'));

    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                       '111111011101101',
                       NULL, 'URI'));
  EXCEPTION
    WHEN others THEN
      dbms_output.put_line('Contractor tag translation failed: '||SQLERRM);
  END;

  -- Test tag translation of employee scheme. Still works.
  BEGIN
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                       'mycompany.employee.123.45',
                       NULL, 'BINARY'));
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                       '011111011101101',
                       NULL, 'URI'));
  EXCEPTION
    WHEN others THEN
      dbms_output.put_line('Employee tag translation failed: '||SQLERRM);
  END;

  -- remove the testing category, which also removes all the associated schemes
  DBMS_MGD_ID_UTL.remove_category('MGD_SAMPLE_CATEGORY', '1.0');
END;
/
SHOW ERRORS;
call DBMS_MGD_ID_UTL.remove_proxy();

@add_scheme3.sql
.
.
.
Validate the MGD_SAMPLE_CATEGORY Scheme
EMPLOYEE_TAG;URI,BINARY;divisionID,employeeID
Length of scheme xml is: 933
111111011101101
mycompany.contractor.123.45
011111011101101
mycompany.employee.123.45
Contractor tag translation failed: ORA-55203: Tag data translation level not found
ORA-06512: at "MGDSYS.DBMS_MGD_ID_UTL", line 54
ORA-06512: at "MGDSYS.MGD_ID", line 242
ORA-29532: Java call terminated by uncaught Java
```

```
exception: oracle.mgd.idcode.exceptions.TDTLevelNotFound: Matching level not
found for any configured scheme
011111011101101
mycompany.employee.123.45
.
.
.
```

Identity Code Package Types

Table 17–2 describes the Identity Code Package ADTs.

Table 17–2 Identity Code Package ADTs

ADT Name	Description
`MGD_ID_COMPONENT` ADT	A data type that specifies the name and value pair attributes that define a component.
`MGD_ID_COMPONENT_VARRAY` ADT	A data type that specifies a list of up to 128 components as name-value attribute pairs used in two constructor functions for creating an identity code type object.
`MGD_ID` ADT	Represents an identity code type that specifies the category identifier for the code category for this identity code and its list of components.

Table 17–3 describes the subprograms in the `MGD_ID` ADT.

All the values and names passed to the subprograms defined in the `MGD_ID` ADT are case-insensitive unless otherwise noted. To preserve case, enclose values in double quotation marks.

Table 17–3 MGD_ID ADT Subprograms

Subprogram	Description
`MGD_ID` Constructor Function	Creates an identity code type object, `MGD_ID`, and returns self.
`FORMAT` Member Function	Returns a representation of an identity code given an `MGD_ID` component.
`GET_COMPONENT` Member Function	Returns the value of an `MGD_ID` component.
`TO_STRING` Member Function	Concatenates the `category_id` parameter value with the components name-value attribute pair.
`TRANSLATE` Static Function	Translates one `MGD_ID` representation of an identity code into a different `MGD_ID` representation.

DBMS_MGD_ID_UTL Package

Table 17–4 describes the Utility subprograms in the `DBMS_MGD_ID_UTL` package.

All the values and names passed to the subprograms defined in the `MGD_ID` ADT are case-insensitive unless otherwise noted. To preserve case, enclose values in double quotation marks.

Table 17–4 DBMS_MGD_ID_UTL Package Utility Subprograms

Subprogram	Description
ADD_SCHEME Procedure	Adds a tag data translation scheme to an existing category.
CREATE_CATEGORY Function	Creates a category or a version of a category.
EPC_TO_ORACLE Function	Converts the EPCglobal tag data translation (TDT) XML to Oracle Database tag data translation XML.
GET_CATEGORY_ID Function	Returns the category ID given the category name and the category version.
GET_COMPONENTS Function	Returns all relevant separated component names separated by semicolon (';') for the specified scheme.
GET_ENCODINGS Function	Returns a list of semicolon (';') separated encodings (formats) for the specified scheme.
GET_JAVA_LOGGING_LEVEL Function	Returns an integer representing the current Java trace logging level.
GET_PLSQL_LOGGING_LEVEL Function	Returns an integer representing the current PL/SQL trace logging level.
GET_SCHEME_NAMES Function	Returns a list of semicolon (';') separated scheme names for the specified category.
GET_TDT_XML Function	Returns the Oracle Database tag data translation XML for the specified scheme.
GET_VALIDATOR Function	Returns the Oracle Database tag data translation schema.
REFRESH_CATEGORY Function	Refreshes the metadata information about the Java stack for the specified category.
REMOVE_CATEORY Function	Removes a category including all the related TDT XML.
REMOVE_PROXY Procedure	Unsets the host and port of the proxy server.
REMOVE_SCHEME Procedure	Removes the tag scheme for a category.
SET_JAVA_LOGGING_LEVEL Procedure	Sets the Java logging level.
SET_PLSQL_LOGGING_LEVEL Procedure	Sets the PL/SQL tracing logging level.
SET_PROXY Procedure	Sets the host and port of the proxy server for Internet access.
VALIDATE_SCHEME Function	Validates the input tag data translation XML against the Oracle Database tag data translation schema.

Identity Code Metadata Tables and Views

This topic describes the structure of identity code metadata tables and views and explains how the metadata are used by the Identity Code Package to interpret the various RFID tags. The creation of these meta tables, views, and triggers is done automatically during the Identity Code Package installation.

Encoding metadata views are used to store encoding categories and schemes. Application developers can insert the meta information of their own identity codes into these views. The MGD_ID ADT is designed to understand the encodings if the metadata for the encodings are stored in the meta tables. If an application developer only uses the encodings defined in the EPC specification v1.1, the developer does not

have to worry about the meta tables because product codes specified in EPC spec v1.1 are predefined.

There are two encoding metadata views.

- `user_mgd_id_category` — this view is used to store the encoding category information defined by the session user.
- `user_mgd_id_scheme` — this view is used to store the encoding type information defined by the session user.

In addition, these read-only views are defined for a user to query the system predefined encoding metadata and the metadata defined by the user:

- `mgd_id_category` — this view is used to query the encoding category information defined by the system or the session user
- `mgd_id_scheme` — this view is used to query the encoding type information defined by the system or the session user.

The underlying metadata tables for the preceding views are:

- `mgd_id_xml_validator`
- `mgd_id_category_tab`
- `mgd_id_scheme_tab`

Users other than the Identity Code Package system users cannot operate on these tables. Users must not use the metadata tables directly. They must use the read only views and the metadata functions described in the `DBMS_MGD_ID_UTL` package.

> **See Also:** *Oracle Database PL/SQL Packages and Types Reference* for information about the `DBMS_MGD_ID_UTL` package

Metadata View Definitions

Table 17–5, Table 17–6, Table 17–7, and Table 17–8 describe the metadata view definitions for the MGD_ID_CATEGORY, USER_ID_CATEGORY, MGD_ID_SCHME, and USER_MGD_ID_SCHME respectively as defined in the `mgdview.sql` file.

Table 17–5 Definition and Description of the MGD_ID_CATEGORY Metadata View

Column Name	Data Type	Description
`CATEGORY_ID`	`NUMBER(4)`	Category identifier
`CATEGORY_NAME`	`VARCHAR2(256)`	Category name
`AGENCY`	`VARCHAR2(256)`	Organization that defined the category
`VERSION`	`VARCHAR2(256)`	Category version
`URI`	`VARCHAR2(256)`	URI that describes the category

Table 17–6 Definition and Description of the USER_MGD_ID_CATEGORY Metadata View

Column Name	Data Type	Description
`CATEGORY_ID`	`NUMBER(4)`	Category identifier
`CATEGORY_NAME`	`VARCHAR2(256)`	Category name
`AGENCY`	`VARCHAR2(256)`	Organization that defined the category
`VERSION`	`VARCHAR2(256)`	Category version
`URI`	`VARCHAR2(256)`	URI that describes the category

Table 17–7 Definition and Description of the MGD_ID_SCHEME Metadata View

Column Name	Data Type	Description
CATEGORY_ID	NUMBER(4)	Category identifier
TYPE_NAME	VARCHAR2(256)	Encoding scheme name, for example, SGTIN-96, GID-96, and so forth
TDT_XML	CLOB	Tag data translation XML for this encoding scheme
ENCODINGS	VARCHAR2(256)	Encodings separated by a comma (,), for example, LEGACY, TAG_ENCODING, PURE_IDENTITY, BINARY (for SGTIN-96)
COMPONENTS	VARCHAR2(1024)	Relevant component names, extracted from each level and then combined. Each is separated by a comma (,). For example, objectclass, generalmanager, serial (for GID-96)

Table 17–8 Definition and Description of the USER_MGD_ID_SCHEME Metadata View

Column Name	Data Type	Description
CATEGORY_ID	NUMBER(4)	Category identifier
TYPE_NAME	VARCHAR2(256)	Encoding scheme name, for example, SGTIN-96, GID-96, and so forth
TDT_XML	CLOB	Tag data translation XML for this encoding scheme
ENCODINGS	VARCHAR2(256)	Encodings separated by a comma (,), for example, LEGACY, TAG_ENCODING, PURE_IDENTITY, BINARY (for SGTIN-96)
COMPONENTS	VARCHAR2(1024)	Relevant component names, extracted from each level and then combined. Each is separated by a comma (,). For example, objectclass, generalmanager, serial (for GID-96)

Electronic Product Code (EPC) Concepts

Topics:

- RFID Technology and EPC v1.1 Coding Schemes
- Product Code Concepts and Their Current Use

RFID Technology and EPC v1.1 Coding Schemes

Radio Frequency Identification (RFID) technology continues to gain momentum with suppliers, distributors, manufacturers, and retailers for its ability to eliminate line-of-site processes and automate critical supply chain transactions. Electronic Product Code (EPC), an identification scheme for universally identifying objects using RFID tags and other means, is gaining widespread acceptance as an emerging standard. Its capabilities enable companies to reduce warehouse and distribution costs through improved inventory control and extended supply chain visibility.

The standardized EPC Identifier is a metacoding scheme designed to support the needs of various industries. Therefore, the EPC represents a family of coding schemes

and a means to make them unique across all possible EPC-compliant tags. EPC Version 1.1 includes these specific coding schemes:

- General Identifier (GID)
- Serialized version of the EAN.UCC Global Trade Item Number (GTIN)
- EAN.UCC Serial Shipping Container Code (SSCC)
- EAN.UCC Global Location Number (GLN)
- EAN.UCC Global Returnable Asset Identifier (GRAI)
- EAN.UCC Global Individual Asset Identifier (GIAI)

RFID applications require the storage of a large volume of EPC data into a database. The efficient use of EPC data also requires that the database recognizes the different coding schemes of EPC data.

EPC is an emerging standard. It does not cover all the numbering schemes used in the various industries and is itself still evolving (the changes from EPC version 1.0 to EPC version 1.1 are significant).

Identity Code Package empowers the Oracle Database with the knowledge to recognize EPC coding schemes. It makes the Oracle Database a database system that not only provides efficient storage and component level retrieval for EPC data, but also has the built-in features to support EPC data encoding and decoding, and conversion between bit encoding and URI encoding.

Identity Code Package provides an extensible framework that allows developers to define their own coding schemes that are not included in the EPC standard. This extensibility feature also makes the Oracle Database adaptable to the evolving future EPC standard.

This chapter describes the requirement of storing, retrieving, encoding and decoding various product codes, including EPC, in an Oracle Database and shows how the Identity Code Package solution meets all these requirements by providing data types, metadata tables, and PL/SQL packages for these purposes.

Product Code Concepts and Their Current Use

This topic describes these product codes:

- Electronic Product Code (EPC)
- Global Trade Identification Number (GTIN) and Serializable Global Trade Identification Number (SGTIN)
- Serial Shipping Container Code (SSCC)
- Global Location Number (GLN) and Serializable Global Location Number (SGLN)
- Global Returnable Asset Identifier (GRAI)
- Global Individual Asset Identifier (GIAI)
- RFID EPC Network

Electronic Product Code (EPC)

The Electronic Product Code™ (EPC™) is an identification scheme for universally identifying physical objects using Radio Frequency Identification (RFID) tags and other means. The standardized EPC data consists of an EPC (or EPC Identifier) that uniquely identifies an individual object, and an optional Filter Value when judged to

be necessary to enable effective and efficient reading of the EPC tags. In addition to this standardized data, certain classes of EPC tags allow user-defined data.

The EPC Identifier is a meta-coding scheme designed to support the needs of various industries by accommodating both existing coding schemes where possible and defining schemes where necessary. The various coding schemes are referred to as Domain Identifiers, to indicate that they provide object identification within certain domains such as a particular industry or group of industries. As such, EPC represents a family of coding schemes (or "namespaces") and a means to make them unique across all possible EPC-compliant tags.

The EPCGlobal EPC Data Standards Version 1.1 defines the abstract content of the Electronic Product Code, and its concrete realization in the form of RFID tags, Internet URIs, and other representations. In EPC Version 1.1, the specific coding schemes include a General Identifier (GID), a serialized version of the EAN.UCC Global Trade Item Number (GTIN®), the EAN.UCC Serial Shipping Container Code (SSCC®), the EAN.UCC Global Location Number (GLN®), the EAN.UCC Global Returnable Asset Identifier (GRAI®), and the EAN.UCC Global Individual Asset Identifier (GIAI®).

EPC Pure Identity The EPC pure identity is the identity associated with a specific physical or logical entity, independent of any particular encoding vehicle such as an RF tag, bar code or database field. As such, a pure identity is an abstract name or number used to identify an entity. A pure identity consists of the information required to uniquely identify a specific entity, and no more.

EPC Encoding EPC encoding is a pure identity with more information, such as filter value, rendered into a specific syntax (typically consisting of value fields of specific sizes). A given pure identity might have several possible encodings, such as a Barcode Encoding, various Tag Encodings, and various URI Encodings. Encodings may also incorporate additional data besides the identity (such as the Filter Value used in some encodings), in which case the encoding scheme specifies what additional data it can hold.

For example, the Serial Shipping Container Code (SSCC) format as defined by the EAN.UCC System is an example of a pure identity. An SSCC encoded into the EPC-SSCC 96-bit format is an example of an encoding.

EPC Tag Bit-Level Encoding EPC encoding on a tag is a string of bits, consisting of a tiered, variable length header followed by a series of numeric fields whose overall length, structure, and function are completely determined by the header value.

EPC Identity URI The EPC identity URI is a representation of a pure identity as a Uniform Resource Identifier (URI).

EPC Tag URI Encoding The EPC tag URI encoding represents a specific EPC tag bit-level encoding, for example, urn:epc:tag:sgtin-64:3.0652642.800031.400.

EPC Encoding Procedure The EPC encoding procedure is used to generate an EPC tag bit-level encoding using various information.

EPC Decoding Procedure The EPC decoding procedure is used to convert an EPC tag bit-level encoding to an EAN.UCC code.

Global Trade Identification Number (GTIN) and Serializable Global Trade Identification Number (SGTIN)

A Global Trade Identification Number (GTIN) is used for the unique identification of trade items worldwide within the EAN.UCC system. The Serialized Global Trade Identification Number (SGTIN) is an identity type in EPC standard version1.1. It is based on the EAN.UCC GTIN code defined in the General EAN.UCC Specifications [GenSpec5.0]. A GTIN identifies a particular class of object, such as a particular kind of product or SKU. The combination of GTIN and a unique serial number is called a Serialized GTIN (SGTIN).

Serial Shipping Container Code (SSCC)

The Serial Shipping Container Code (SSCC) is defined by the General EAN.UCC Specifications [GenSpec5.0]. The unique identification of logistics units is achieved in the EAN.UCC system by the use of the SSCC. The SSCC is intended for assignment to individual objects.

Global Location Number (GLN) and Serializable Global Location Number (SGLN)

The Global Location Number (GLN) is defined by the General EAN.UCC Specifications [GenSpec5.0]. A GLN can represent either a discrete, unique physical location such as a dock door or a warehouse slot, or an aggregate physical location such as an entire warehouse. In addition, a GLN can represent a logical entity such as an organization that performs a business function (for example, placing an order). The combination of GLN and a unique serial number is called a Serialized GLN (SGLN). However, until the EAN.UCC community determines the appropriate way to extend GLN, the serial number field is reserved and must not be used.

Global Returnable Asset Identifier (GRAI)

A returnable asset is a reusable package or transport equipment of a certain value. Global Returnable Asset Identifier is (GRAI) is defined by the General EAN.UCC Specifications [GenSpec5.0] for the unique identification of a returnable asset.

Global Individual Asset Identifier (GIAI)

The Global Individual Asset Identifier (GIAI) is defined by the General EAN.UCC Specifications [GenSpec5.0]. Unlike the GTIN, the GIAI is intended for assignment to individual objects. Global Individual Asset Identifier (GIAI) is used to uniquely identify an entity that is part of the fixed inventory of a company. The GIAI identifies any fixed asset of an organization.

RFID EPC Network

The RFID EPC network is used to identify, track and locate assets. Physical objects are identified by a unique RFID enabled EPC.

Oracle Database Tag Data Translation Schema

The Oracle Database Tag Data Translation Schema is closely related to the EPCglobal TDT schema, however it is not exact. The Oracle Database TDT is shown as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema targetNamespace="oracle.mgd.idcode"
   xmlns:xsd="http://www.w3.org/2001/XMLSchema"
     xmlns:tdt="oracle.mgd.idcode" elementFormDefault="qualified"
       attributeFormDefault="unqualified" version="1.0">
```

```
<xsd:simpleType name="InputFormatList">
 <xsd:restriction base="xsd:string">
  <xsd:enumeration value="BINARY"/>
  <xsd:enumeration value="STRING"/>
 </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="LevelTypeList">
 <xsd:restriction base="xsd:string">
 </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="SchemeNameList">
 <xsd:restriction base="xsd:string">
 </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="ModeList">
 <xsd:restriction base="xsd:string">
  <xsd:enumeration value="EXTRACT"/>
  <xsd:enumeration value="FORMAT"/>
 </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="CompactionMethodList">
 <xsd:restriction base="xsd:string">
  <xsd:enumeration value="32-bit"/>
  <xsd:enumeration value="16-bit"/>
  <xsd:enumeration value="8-bit"/>
  <xsd:enumeration value="7-bit"/>
  <xsd:enumeration value="6-bit"/>
  <xsd:enumeration value="5-bit"/>
 </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="PadDirectionList">
 <xsd:restriction base="xsd:string">
  <xsd:enumeration value="LEFT"/>
  <xsd:enumeration value="RIGHT"/>
 </xsd:restriction>
</xsd:simpleType>

<xsd:complexType name="Field">
 <xsd:attribute name="seq" type="xsd:integer" use="required"/>
 <xsd:attribute name="name" type="xsd:string" use="required"/>
 <xsd:attribute name="bitLength" type="xsd:integer"/>
 <xsd:attribute name="characterSet" type="xsd:string" use="required"/>
 <xsd:attribute name="compaction" type="tdt:CompactionMethodList"/>
 <xsd:attribute name="compression" type="xsd:string"/>
 <xsd:attribute name="padChar" type="xsd:string"/>
 <xsd:attribute name="padDir" type="tdt:PadDirectionList"/>
 <xsd:attribute name="decimalMinimum" type="xsd:long"/>
 <xsd:attribute name="decimalMaximum" type="xsd:long"/>
 <xsd:attribute name="length" type="xsd:integer"/>
</xsd:complexType>

<xsd:complexType name="Option">
 <xsd:sequence>
  <xsd:element name="field" type="tdt:Field" maxOccurs="unbounded"/>
 </xsd:sequence>
 <xsd:attribute name="optionKey" type="xsd:string" use="required"/>
```

```
  <xsd:attribute name="pattern" type="xsd:string"/>
  <xsd:attribute name="grammar" type="xsd:string" use="required"/>
 </xsd:complexType>

 <xsd:complexType name="Rule">
  <xsd:attribute name="type" type="tdt:ModeList" use="required"/>
  <xsd:attribute name="inputFormat" type="tdt:InputFormatList" use="required"/>
  <xsd:attribute name="seq" type="xsd:integer" use="required"/>
  <xsd:attribute name="newFieldName" type="xsd:string" use="required"/>
  <xsd:attribute name="characterSet" type="xsd:string" use="required"/>
  <xsd:attribute name="padChar" type="xsd:string"/>
  <xsd:attribute name="padDir" type="tdt:PadDirectionList"/>
  <xsd:attribute name="decimalMinimum" type="xsd:long"/>
  <xsd:attribute name="decimalMaximum" type="xsd:long"/>
  <xsd:attribute name="length" type="xsd:string"/>
  <xsd:attribute name="function" type="xsd:string" use="required"/>
  <xsd:attribute name="tableURI" type="xsd:string"/>
  <xsd:attribute name="tableParams" type="xsd:string"/>
  <xsd:attribute name="tableXPath" type="xsd:string"/>
  <xsd:attribute name="tableSQL" type="xsd:string"/>
 </xsd:complexType>

 <xsd:complexType name="Level">
  <xsd:sequence>
   <xsd:element name="option" type="tdt:Option" minOccurs="1"
     maxOccurs="unbounded"/>
   <xsd:element name="rule" type="tdt:Rule" minOccurs="0"
     maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="type" type="tdt:LevelTypeList" use="required"/>
  <xsd:attribute name="prefixMatch" type="xsd:string"/>
  <xsd:attribute name="requiredParsingParameters" type="xsd:string"/>
  <xsd:attribute name="requiredFormattingParameters" type="xsd:string"/>
 </xsd:complexType>

 <xsd:complexType name="Scheme">
  <xsd:sequence>
   <xsd:element name="level" type="tdt:Level" minOccurs="4" maxOccurs="5"/>
  </xsd:sequence>
  <xsd:attribute name="name" type="tdt:SchemeNameList" use="required"/>
  <xsd:attribute name="optionKey" type="xsd:string" use="required"/>
 </xsd:complexType>
 <xsd:complexType name="TagDataTranslation">
  <xsd:sequence>
   <xsd:element name="scheme" type="tdt:Scheme" maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="version" type="xsd:string" use="required"/>
  <xsd:attribute name="date" type="xsd:dateTime" use="required"/>
 </xsd:complexType>
 <xsd:element name="TagDataTranslation" type="tdt:TagDataTranslation"/>
</xsd:schema>
```

18

Schema Object Dependency

If the definition of object A references object B, then A depends on B. This chapter explains dependencies among schema objects, and how Oracle Database automatically tracks and manages these dependencies. Because of this automatic dependency management, A never uses an obsolete version of B, and you almost never have to explicitly recompile A after you change B.

Topics:

- Overview of Schema Object Dependencies
- Querying Object Dependencies
- Object Status
- Invalidation of Dependent Objects
- Guidelines for Reducing Invalidation
- Object Revalidation
- Name Resolution in Schema Scope
- Local Dependency Management
- Remote Dependency Management
- Remote Procedure Call (RPC) Dependency Management
- Shared SQL Dependency Management

Overview of Schema Object Dependencies

Some types of schema objects can reference other objects in their definitions. For example, a view is defined by a query that references tables or other views, and the body of a subprogram can include SQL statements that reference other objects. If the definition of object A references object B, then A is a **dependent object** (of B) and B is a **referenced object** (of A).

Example 18–1 shows how to display the dependent and referenced object types in your database (if you are logged in as DBA).

Example 18–1 Displaying Dependent and Referenced Object Types

Display dependent object types:

```
SELECT DISTINCT TYPE
FROM DBA_DEPENDENCIES
ORDER BY TYPE;
```

Result:

```
TYPE
------------------
DIMENSION
EVALUATION CONTXT
FUNCTION
INDEX
INDEXTYPE
JAVA CLASS
JAVA DATA
MATERIALIZED VIEW
OPERATOR
PACKAGE
PACKAGE BODY

TYPE
------------------
PROCEDURE
RULE
RULE SET
SYNONYM
TABLE
TRIGGER
TYPE
TYPE BODY
UNDEFINED
VIEW
XML SCHEMA

22 rows selected.
```

Display referenced object types:

```
SELECT DISTINCT REFERENCED_TYPE
FROM DBA_DEPENDENCIES
ORDER BY REFERENCED_TYPE;
```

Result:

```
REFERENCED_TYPE
------------------
EVALUATION CONTXT
FUNCTION
INDEX
INDEXTYPE
JAVA CLASS
LIBRARY
NON-EXISTENT
OPERATOR
PACKAGE
PROCEDURE
SEQUENCE

REFERENCED_TYPE
------------------
SYNONYM
TABLE
TYPE
VIEW
XML SCHEMA
```

```
16 rows selected.
```

If you alter the definition of a referenced object, dependent objects might not continue to function without error, depending on the type of alteration. For example, if you drop a table, no view based on the dropped table is usable.

As an example of a schema object change that invalidates some dependents but not others, consider the two views in Example 18–2, which are based on the HR.EMPLOYEES table.

Example 18–2 creates two views from the EMPLOYEES table: SIXFIGURES, which selects all columns in the table, and COMMISSIONED, which does not include the EMAIL column. As the example shows, changing the EMAIL column invalidates SIXFIGURES, but not COMMISSIONED.

Example 18–2 Schema Object Change that Invalidates Some Dependents

```
CREATE OR REPLACE VIEW sixfigures AS
SELECT * FROM employees
WHERE salary >= 100000;

CREATE OR REPLACE VIEW commissioned AS
SELECT first_name, last_name, commission_pct
FROM employees
WHERE commission_pct > 0.00;
```

SQL*Plus formatting command:

```
COLUMN object_name FORMAT A16
```

Query:

```
SELECT object_name, status
FROM user_objects
WHERE object_type = 'VIEW'
ORDER BY object_name;
```

Result:

```
OBJECT_NAME      STATUS
---------------- -------
COMMISSIONED     VALID
SIXFIGURES       VALID
```

Lengthen EMAIL column of EMPLOYEES table:

```
ALTER TABLE employees MODIFY email VARCHAR2(100);
```

Query:

```
SELECT object_name, status
FROM user_objects
WHERE object_type = 'VIEW'
ORDER BY object_name;
```

Result:

```
OBJECT_NAME      STATUS
---------------- -------
COMMISSIONED     INVALID
SIXFIGURES       VALID
```

A view depends on every object referenced in its query. The view in Example 18–3, depends on the tables `employees` and `departments`.

Example 18–3 View that Depends on Multiple Objects

```
CREATE OR REPLACE VIEW v AS
  SELECT last_name, first_name, department_name
  FROM employees e, departments d
  WHERE e.department_id = d.department_id
  ORDER BY last_name;
```

Notes:

- `CREATE` statements automatically update all dependencies.
- Dynamic SQL statements do not create dependencies. For example, this statement does not create a dependency on `tab1`:

  ```
  EXECUTE IMMEDIATE 'SELECT * FROM tab1'
  ```

Querying Object Dependencies

The static data dictionary views `USER_DEPENDENCIES`, `ALL_DEPENDENCIES`, and `DBA_DEPENDENCIES` describe dependencies between database objects.

The `utldtree.sql` SQL script creates the view `DEPTREE`, which contains information on the object dependency tree, and the view `IDEPTREE`, a presorted, pretty-print version of `DEPTREE`.

See Also: *Oracle Database Reference* for more information about the `DEPTREE`, `IDEPTREE`, and `utldtree.sql` script

Object Status

Every database object has a status value described in Table 18–1.

Table 18–1 Database Object Status

Status	Meaning
Valid	The object was successfully compiled, using the current definition in the data dictionary.
Compiled with errors	The most recent attempt to compile the object produced errors.
Invalid	The object is marked invalid because an object that it references has changed. (Only a dependent object can be invalid.)
Unauthorized	An access privilege on a referenced object was revoked. (Only a dependent object can be unauthorized.)

Note: The static data dictionary views `USER_OBJECTS`, `ALL_OBJECTS`, and `DBA_OBJECTS` do not distinguish between "Compiled with errors," "Invalid," and "Unauthorized"—they describe all of these as `INVALID`.

Invalidation of Dependent Objects

If object A depends on object B, which depends on object C, then A is a **direct dependent** of B, B is a direct dependent of C, and A is an **indirect dependent** of C.

Direct dependents are invalidated only by changes to the referenced object that affect them (changes to the signature of the referenced object).

Indirect dependents can be invalidated by changes to the reference object that do not affect them. If a change to C invalidates B, it invalidates A (and all other direct and indirect dependents of B). This is called **cascading invalidation**.

With **coarse-grained invalidation**, a data definition language (DDL) statement that changes a referenced object invalidates all of its dependents.

With **fine-grained invalidation**, a DDL statement that changes a referenced object invalidates only dependents for which either of these statements is true:

- The dependent relies on the attribute of the referenced object that the DDL statement changed.
- The compiled metadata of the dependent is no longer correct for the changed referenced object.

For example, if view `v` selects columns `c1` and `c2` from table `t`, a DDL statement that changes only column `c3` of `t` does not invalidate `v`.

The DDL statement `CREATE OR REPLACE` *`object`* has no effect under these conditions:

- *`object`* is a PL/SQL object, the new PL/SQL source text is identical to the existing PL/SQL source text, and the PL/SQL compilation parameter settings stored with *`object`* are identical to those in the session environment.

  For information about PL/SQL compilation parameter settings, see *Oracle Database PL/SQL Language Reference*.
- *`object`* is a synonym and the statement does not change the target object.

The operations in the left column of Table 18–2 cause fine-grained invalidation, except in the cases in the right column. The cases in the right column, and all operations not listed in Table 18–2, cause coarse-grained invalidation.

Table 18–2 Operations that Cause Fine-Grained Invalidation

Operation	Exceptions
`ALTER TABLE` *`table`* `ADD` *`column`*	▪ Dependent object (except a view) uses `SELECT *` on *`table`*. ▪ Dependent object uses *`table`*`%rowtype`. ▪ Dependent object performs `INSERT` on *`table`* without specifying column list. ▪ Dependent object references *`table`* in query that contains a join. ▪ Dependent object references *`table`* in query that references a PL/SQL variable.

Table 18–2 (Cont.) Operations that Cause Fine-Grained Invalidation

Operation	Exceptions
`ALTER TABLE` *`table`* {`MODIFY` | `RENAME` | `DROP` | `SET UNUSED`} *`column`* `ALTER TABLE` *`table`* `DROP CONSTRAINT` *`not_null_constraint`*	■ Dependent object directly references *`column`*. ■ Dependent object uses `SELECT *` on *`table`*. ■ Dependent object uses *`table`*`%ROWTYPE`. ■ Dependent object performs `INSERT` on *`table`* without specifying column list. ■ Dependent object is a trigger that depends on an entire row (that is, it does not specify a column in its definition). ■ Dependent object is a trigger that depends on a column to the right of the dropped column.
`CREATE OR REPLACE VIEW` *`view`* Online Table Redefinition (`DBMS_REDEFINITION`)	Column lists of new and old definitions differ, and at least one of these is true: ■ Dependent object references column that is modified or dropped in new view or table definition. ■ Dependent object uses *`view`*`%rowtype` or *`table`*`%rowtype`. ■ Dependent object performs `INSERT` on view or table without specifying column list. ■ New view definition introduces new columns, and dependent object references view or table in query that contains a join. ■ New view definition introduces new columns, and dependent object references view or table in query that references a PL/SQL variable. ■ Dependent object references view or table in `RELIES ON` clause.
`CREATE OR REPLACE SYNONYM` *`synonym`*	■ New and old *`synonym`* targets differ, and one is not a table. ■ Both old and new *`synonym`* targets are tables, and the tables have different column lists or different privilege grants. ■ Both old and new *`synonym`* targets are tables, and dependent object is a view that references a column that participates in a unique index on the old target but not in a unique index on the new target.

Table 18–2 (Cont.) Operations that Cause Fine-Grained Invalidation

Operation	Exceptions
`DROP INDEX`	■ The index is a function-based index and the dependent object is a trigger that depends either on an entire row or on a column that was added to *`table`* after a function-based index was created. ■ The index is a unique index, the dependent object is a view, and the view references a column participating in the unique index.
`CREATE OR REPLACE {PROCEDURE|FUNCTION}`	Call signature changes. Call signature is the parameter list (order, names, and types of parameters), return type, purity[1], determinism, parallelism, pipelining, and (if the procedure or function is implemented in C or Java) implementation properties.
`CREATE OR REPLACE PACKAGE`	■ Dependent object references a dropped or renamed package item. ■ Dependent object references a package procedure or function whose call signature or entry-point number[2], changed. If referenced procedure or function has multiple overload candidates, dependent object is invalidated if any overload candidate's call signature or entry point number changed, or if a candidate was added or dropped. ■ Dependent object references a package cursor whose call signature, rowtype, or entry point number changed. ■ Dependent object references a package type or subtype whose definition changed. ■ Dependent object references a package variable or constant whose name, data type, initial value, or offset number changed. ■ Package purity[1] changed.

[1] **Purity** refers to a set of rules for preventing side effects (such as unexpected data changes) when invoking PL/SQL functions within SQL queries. **Package purity** refers to the purity of the code in the package initialization block.

[2] The **entry-point number** of a procedure or function is determined by its location in the PL/SQL package code. A procedure or function added to the end of a PL/SQL package is given a new entry-point number.

Note: A dependent object that is invalidated by an operation in Table 18–2 appears in the static data dictionary views `*_OBJECTS` and `*_OBJECTS_AE` only after an attempt to reference it (either during compilation or execution) or after invoking one of these subprograms:

- `DBMS_UTILITY.COMPILE_SCHEMA` (described in *Oracle Database PL/SQL Packages and Types Reference*)
- Any `UTL_RECOMP` subprogram (described in *Oracle Database PL/SQL Packages and Types Reference*)

Topics:

- Session State and Referenced Packages
- Security Authorization

Session State and Referenced Packages

Each session that references a package construct has its own instantiation of that package, including a persistent state of any public and private variables, cursors, and constants. All of a session's package instantiations, including state, can be lost if any of the session's instantiated packages are subsequently invalidated and revalidated.

See Also:

- *Oracle Database PL/SQL Language Reference* for information about package instantiation
- *Oracle Database PL/SQL Language Reference* for information about package state

Security Authorization

When a data manipulation language (DML) object or system privilege is granted to, or revoked from, a user or `PUBLIC`, Oracle Database invalidates all the owner's dependent objects, to verify that an owner of a dependent object continues to have the necessary privileges for all referenced objects.

Guidelines for Reducing Invalidation

To reduce invalidation of dependent objects, follow these guidelines:

- Add Items to End of Package
- Reference Each Table Through a View

Add Items to End of Package

When adding items to a package, add them to the end of the package. This preserves the entry point numbers of existing top-level package items, preventing their invalidation.

For example, consider this package:

```
CREATE OR REPLACE PACKAGE pkg1 IS
  FUNCTION get_var RETURN VARCHAR2;
END;
/
```

Adding an item to the end of pkg1, as follows, does not invalidate dependents that reference the get_var function:

```
CREATE OR REPLACE PACKAGE pkg1 IS
  FUNCTION get_var RETURN VARCHAR2;
  PROCEDURE set_var (v VARCHAR2);
END;
/
```

Inserting an item between the get_var function and the set_var procedure, as follows, invalidates dependents that reference the set_var function:

```
CREATE OR REPLACE PACKAGE pkg1 IS
  FUNCTION get_var RETURN VARCHAR2;
  PROCEDURE assert_var (v VARCHAR2);
  PROCEDURE set_var (v VARCHAR2);
END;
/
```

Reference Each Table Through a View

Reference tables indirectly, using views, enabling you to:

- Add columns to the table without invalidating dependent views or dependent PL/SQL objects
- Modify or delete columns not referenced by the view without invalidating dependent objects

The statement CREATE OR REPLACE VIEW does not invalidate an existing view or its dependents if the new ROWTYPE matches the old ROWTYPE.

Object Revalidation

An object that is not valid when it is referenced must be validated before it can be used. Validation occurs automatically when an object is referenced; it does not require explicit user action.

If an object is not valid, its status is either compiled with errors, unauthorized, or invalid. For definitions of these terms, see Table 18–1.

Topics:

- Revalidation of Objects that Compiled with Errors
- Revalidation of Unauthorized Objects
- Revalidation of Invalid SQL Objects
- Revalidation of Invalid PL/SQL Objects

Revalidation of Objects that Compiled with Errors

The compiler cannot automatically revalidate an object that compiled with errors. The compiler recompiles the object, and if it recompiles without errors, it is revalidated; otherwise, it remains invalid.

Revalidation of Unauthorized Objects

The compiler checks whether the unauthorized object has access privileges to all of its referenced objects. If so, the compiler revalidates the unauthorized object without recompiling it. If not, the compiler issues appropriate error messages.

Revalidation of Invalid SQL Objects

The SQL compiler recompiles the invalid object. If the object recompiles without errors, it is revalidated; otherwise, it remains invalid.

Revalidation of Invalid PL/SQL Objects

For an invalid PL/SQL program unit (procedure, function, or package), the PL/SQL compiler checks whether any referenced object changed in a way that affects the invalid object. If so, the compiler recompiles the invalid object. If the object recompiles without errors, it is revalidated; otherwise, it remains invalid. If not, the compiler revalidates the invalid object without recompiling it.

Name Resolution in Schema Scope

Object names referenced in SQL statements have one or more pieces. Pieces are separated by periods—for example, `hr.employees.department_id` has three pieces.

Oracle Database uses this procedure to try to resolve an object name:

1. Try to qualify the first piece of the object name.

   If the object name has only one piece, then that piece is the first piece. Otherwise, the first piece is the piece to the left of the leftmost period; for example, in `hr.employees.department_id`, the first piece is `hr`.

   The procedure for trying to qualify the first piece is:

   a. If the object name is a table name that appears in the `FROM` clause of a `SELECT` statement, and the object name has multiple pieces, go to step d. Otherwise, go to step b.

   b. Search the current schema for an object whose name matches the first piece.

      If found, go to step 2. Otherwise, go to step c.

   c. Search for a public synonym that matches the first piece.

      If found, go to step 2. Otherwise, go to step d.

   d. Search for a schema whose name matches the first piece.

      If found, and if the object name has a second piece, go to step e. Otherwise, return an error—the object name cannot be qualified.

   e. Search the schema found at step d for a built-in function whose name matches the second piece of the object name.

      If found, the schema redefined that built-in function. The object name resolves to the original built-in function, not to the schema-defined function of the same name. Go to step 2.

      If not found, return an error—the object name cannot be qualified.

2. A schema object has been qualified. Any remaining pieces of the object name must match a valid part of this schema object.

   For example, if the object name is `hr.employees.department_id`, `hr` is qualified as a schema. If `employees` is qualified as a table, `department_id` must

correspond to a column of that table. If `employees` is qualified as a package, `department_id` must correspond to a public constant, variable, procedure, or function of that package.

Because of how Oracle Database resolves references, an object can depend on the nonexistence of other objects. This situation occurs when the dependent object uses a reference that would be interpreted differently if another object were present.

See Also:

- *Oracle Database PL/SQL Language Reference* for information about how name resolution differs in SQL and PL/SQL
- *Oracle Database Administrator's Guide* for information about name resolution in a distributed database system

Local Dependency Management

Local dependency management occurs when Oracle Database manages dependencies among the objects in a single database. For example, a statement in a procedure can reference a table in the same database.

Remote Dependency Management

Remote dependency management occurs when Oracle Database manages dependencies in distributed environments across a network. For example, an Oracle Forms trigger can depend on a schema object in the database. In a distributed database, a local view can reference a remote table.

Oracle Database also manages distributed database dependencies. For example, an Oracle Forms application might contain a trigger that references a table. The database system must account for dependencies among such objects. Oracle Database uses different mechanisms to manage remote dependencies, depending on the objects involved.

Topics:

- Dependencies Among Local and Remote Database Procedures
- Dependencies Among Other Remote Objects
- Dependencies of Applications

Dependencies Among Local and Remote Database Procedures

Dependencies among stored procedures (including functions, packages, and triggers) in a distributed database system are managed using either time-stamp checking or signature checking (see "Time-Stamp Dependency Mode" on page 18-12 and "RPC-Signature Dependency Mode" on page 18-13).

The dynamic initialization parameter `REMOTE_DEPENDENCIES_MODE` determines whether time stamps or signatures govern remote dependencies.

See Also: *Oracle Database PL/SQL Language Reference*

Dependencies Among Other Remote Objects

Oracle Database does not manage dependencies among remote schema objects other than local-procedure-to-remote-procedure dependencies.

For example, assume that a local view is created and defined by a query that references a remote table. Also assume that a local procedure includes a SQL statement that references the same remote table. Later, the definition of the table is altered.

Therefore, the local view and procedure are never invalidated, even if the view or procedure is used after the table is altered, and even if the view or procedure now returns errors when used. In this case, the view or procedure must be altered manually so that errors are not returned. In such cases, lack of dependency management is preferable to unnecessary recompilations of dependent objects.

Dependencies of Applications

Code in database applications can reference objects in the connected database. For example, Oracle Call Interface (OCI) and precompiler applications can submit anonymous PL/SQL blocks. Triggers in Oracle Forms applications can reference a schema object.

Such applications are dependent on the schema objects they reference. Dependency management techniques vary, depending on the development environment. Oracle Database does not automatically track application dependencies.

> **See Also:** Manuals for your application development tools and your operating system for more information about managing the remote dependencies within database applications

Remote Procedure Call (RPC) Dependency Management

Remote procedure call (RPC) dependency management occurs when a local stored procedure calls a remote procedure in a distributed database system. The dynamic initialization parameter `REMOTE_DEPENDENCIES_MODE` controls the dependency mode. The choice is either time-stamp dependency mode or RPC-signature dependency mode.

Topics:

- Time-Stamp Dependency Mode
- RPC-Signature Dependency Mode
- Controlling Dependency Mode

Time-Stamp Dependency Mode

Whenever a procedure is compiled, its **time stamp** is recorded in the data dictionary. The time stamp shows when the procedure was created, altered, or replaced.

A compiled procedure contains information about each remote procedure that it calls, including the schema, package name, procedure name, and time stamp of the remote procedure.

In time-stamp dependency mode, when a local stored procedure calls a remote procedure, Oracle Database compares the time stamp that the local procedure has for the remote procedure to the current time stamp of the remote procedure. If the two timestamps match, both the local and remote procedures run. Neither is recompiled.

If the two timestamps do not match, the local procedure is invalidated and an error is returned to the calling environment. All other local procedures that depend on the remote procedure with the new time stamp are also invalidated.

Time stamp comparison occurs when a statement in the body of the local procedure calls the remote procedure. Therefore, statements in the local procedure that precede the invalid call might run successfully. Statements after the invalid call do not run. The local procedure must be recompiled.

If DML statements precede the invalid call, they roll back only if they and the invalid call are in the same PL/SQL block. For example, the `UPDATE` statement rolls back in this code:

```
BEGIN
  UPDATE table SET ...
  invalid_proc;
  COMMIT;
END;
```

But the `UPDATE` statement does not roll back in this code:

```
UPDATE table SET ...
EXECUTE invalid_proc;
COMMIT;
```

The disadvantages of time-stamp dependency mode are:

- Dependent objects across the network are often recompiled unnecessarily, degrading performance.
- If the client-side application uses PL/SQL version 2, this mode can cause situations that prevent the application from running on the client side.

  An example of such an application is any release of Oracle Forms that is integrated with PL/SQL version 2 on the client side. During installation, you must recompile the client-side PL/SQL procedures that Oracle Forms uses at the client site. Also, if a client-side procedure depends on a server procedure, and if the server procedure changes or is automatically recompiled, you must recompile the client-side PL/SQL procedure. However, no PL/SQL compiler is available on the client. Therefore, the developer of the client application must distribute new versions of the application to all customers.

  Client-side applications that used PL/SQL version 1, such as earlier releases of Oracle Forms, did not use time-stamp dependency mode, because PL/SQL version 1 did not support stored procedures.

RPC-Signature Dependency Mode

Oracle Database provides **RPC signatures** to handle remote dependencies. RPC signatures do not affect local dependencies, because recompilation is always possible in the local environment.

An RPC signature is associated with each compiled stored program unit. It identifies the unit by these characteristics:

- Name
- Number of parameters
- Data type class of each parameter
- Mode of each parameter
- Data type class of return value (for a function)

An RPC signature changes only when at least one of the preceding characteristics changes.

> **Note:** An RPC signature does not include DETERMINISTIC, PARALLEL_ENABLE, or purity information. If these settings change for a function on remote system, optimizations based on them are not automatically reconsidered. Therefore, calling the remote function in a SQL statement or using it in a function-based index might cause incorrect query results.

A compiled program unit contains the RPC signature of each remote procedure that it calls (and the schema, package name, procedure name, and time stamp of the remote procedure).

In RPC-signature dependency mode, when a local program unit calls a subprogram in a remote program unit, the database ignores time-stamp mismatches and compares the RPC signature that the local unit has for the remote subprogram to the current RPC signature of the remote subprogram. If the RPC signatures match, the call succeeds; otherwise, the database returns an error to the local unit, and the local unit is invalidated.

For example, suppose that this procedure, get_emp_name, is stored on a server in Boston (BOSTON_SERVER):

```
CREATE OR REPLACE PROCEDURE get_emp_name (
  emp_number  IN  NUMBER,
  hiredate    OUT VARCHAR2,
  emp_name    OUT VARCHAR2) AS
BEGIN
  SELECT last_name, TO_CHAR(hire_date, 'DD-MON-YY')
  INTO emp_name, hiredate
  FROM employees
  WHERE employee_id = emp_number;
END;
/
```

When get_emp_name is compiled on BOSTON_SERVER, Oracle Database records both its RPC signature and its time stamp.

Suppose that this PL/SQL procedure, print_name, which calls get_emp_name, is on a server in California:

```
CREATE OR REPLACE PROCEDURE print_ename (emp_number IN NUMBER) AS
  hiredate  VARCHAR2(12);
  ename     VARCHAR2(10);
BEGIN
  get_emp_name@BOSTON_SERVER(emp_number, hiredate, ename);
  dbms_output.put_line(ename);
  dbms_output.put_line(hiredate);
END;
/
```

When print_name is compiled on the California server, the database connects to the Boston server, sends the RPC signature of get_emp_name to the California server, and records the RPC signature of get_emp_name in the compiled state of print_ename.

At run time, when print_name calls get_emp_name, the database sends the RPC signature of get_emp_name that was recorded in the compiled state of print_ename to the Boston server. If the recorded RPC signature matches the current RPC signature of get_emp_name on the Boston server, the call succeeds; otherwise, the database returns an error to print_name, which is invalidated.

Topics:

- Changing Names and Default Values of Parameters
- Changing Specification of Parameter Mode IN
- Changing Subprogram Body
- Changing Data Type Classes of Parameters
- Changing Package Types

Changing Names and Default Values of Parameters

Changing the name or default value of a subprogram parameter does not change the RPC signature of the subprogram. For example, procedure `P1` has the same RPC signature in these two examples:

```
PROCEDURE P1 (Param1 IN NUMBER := 100);
PROCEDURE P1 (Param2 IN NUMBER := 200);
```

However, if your application requires that callers get the new default value, you must recompile the called procedure.

Changing Specification of Parameter Mode IN

Because the subprogram parameter mode `IN` is the default, you can specify it either implicitly or explicitly. Changing its specification from implicit to explicit, or the reverse, does not change the RPC signature of the subprogram. For example, procedure `P1` has the same RPC signature in these two examples:

```
PROCEDURE P1 (Param1 NUMBER);    -- implicit specification
PROCEDURE P1 (Param1 IN NUMBER);  -- explicit specification
```

Changing Subprogram Body

Changing the body of a subprogram does not change the RPC signature of the subprogram.

Example 18–4 changes only the body of the procedure `get_hire_date`; therefore, it does not change the RPC signature of `get_hire_date`.

Example 18–4 Changing Body of Procedure get_hire_date

```
CREATE OR REPLACE PROCEDURE get_hire_date (
  emp_number  IN  NUMBER,
  hiredate    OUT VARCHAR2,
  emp_name    OUT VARCHAR2) AS
BEGIN
  SELECT last_name, TO_CHAR(hire_date, 'DD-MON-YY')
  INTO emp_name, hiredate
  FROM employees
  WHERE employee_id = emp_number;
END;
/

CREATE OR REPLACE PROCEDURE get_hire_date (
  emp_number  IN  NUMBER,
  hiredate    OUT VARCHAR2,
  emp_name    OUT VARCHAR2) AS
BEGIN
  -- Change date format model
  SELECT last_name, TO_CHAR(hire_date, 'DD/MON/YYYY')
```

```
  INTO emp_name, hiredate
  FROM employees
  WHERE employee_id = emp_number;
END;
/
```

Changing Data Type Classes of Parameters

Changing the data type of a parameter to another data type in the same class does not change the RPC signature, but changing the data type to a data type in another class does.

Table 18–3 lists the data type classes and the data types that comprise them. Data types not listed in Table 18–3, such as NCHAR, do not belong to a data type class. Changing their type always changes the RPC signature.

Table 18–3 Data Type Classes

Data Type Class	Data Types in Class
Character	CHAR CHARACTER
VARCHAR	VARCHAR VARCHAR2 STRING LONG ROWID
Raw	RAW LONG RAW
Integer	BINARY_INTEGER PLS_INTEGER SIMPLE_INTEGER BOOLEAN NATURAL NATURALN POSITIVE POSITIVEN
Number	NUMBER INT INTEGER SMALLINT DEC DECIMAL REAL FLOAT NUMERIC DOUBLE PRECISION
Date	DATE TIMESTAMP TIMESTAMP WITH TIME ZONE TIMESTAMP WITH LOCAL TIME ZONE INTERVAL YEAR TO MONTH INTERVAL DAY TO SECOND

Example 18–5 changes the data type of the parameter hiredate from VARCHAR2 to DATE. VARCHAR2 and DATE are not in the same data type class, so the RPC signature of the procedure get_hire_date changes.

Example 18–5 Changing Data Type Class of get_hire_date Parameter

```
CREATE OR REPLACE PROCEDURE get_hire_date (
  emp_number  IN  NUMBER,
  hiredate    OUT DATE,
  emp_name    OUT VARCHAR2) AS
BEGIN
  SELECT last_name, TO_CHAR(hire_date, 'DD/MON/YYYY')
  INTO emp_name, hiredate
  FROM employees
  WHERE employee_id = emp_number;
END;
/
```

Changing Package Types

Changing the name of a package type, or the names of its internal components, does not change the RPC signature of the package.

Example 18–6 defines a record type, `emp_data_type`, inside the package `emp_package`. Next, it changes the names of the record fields, but not their types. Finally, it changes the name of the type, but not its characteristics. The RPC signature of the package does not change.

Example 18–6 Changing Names of Fields in Package Record Type

```
CREATE OR REPLACE PACKAGE emp_package AS
  TYPE emp_data_type IS RECORD (
    emp_number  NUMBER,
    hiredate    VARCHAR2(12),
    emp_name    VARCHAR2(10)
  );
  PROCEDURE get_emp_data (
    emp_data  IN OUT emp_data_type
  );
END;
/

CREATE OR REPLACE PACKAGE emp_package AS
  TYPE emp_data_type IS RECORD (
    emp_num   NUMBER,
    hire_dat  VARCHAR2(12),
    empname   VARCHAR2(10)
  );
  PROCEDURE get_emp_data (
    emp_data  IN OUT emp_data_type
  );
END;
/

CREATE OR REPLACE PACKAGE emp_package AS
  TYPE emp_data_record_type IS RECORD (
    emp_num   NUMBER,
    hire_dat  VARCHAR2(12),
    empname   VARCHAR2(10)
  );
  PROCEDURE get_emp_data (
    emp_data  IN OUT emp_data_record_type
  );
END;
/
```

Controlling Dependency Mode

The dynamic initialization parameter REMOTE_DEPENDENCIES_MODE controls the dependency mode. If the initialization parameter file contains this specification, then only time stamps are used to resolve dependencies (if this is not explicitly overridden dynamically):

```
REMOTE_DEPENDENCIES_MODE = TIMESTAMP
```

If the initialization parameter file contains this parameter specification, then RPC signatures are used to resolve dependencies (if this not explicitly overridden dynamically):

```
REMOTE_DEPENDENCIES_MODE = SIGNATURE
```

You can alter the mode dynamically by using the DDL statements. For example, this example alters the dependency mode for the current session:

```
ALTER SESSION SET REMOTE_DEPENDENCIES_MODE = {SIGNATURE | TIMESTAMP}
```

This example alters the dependency mode systemwide after startup:

```
ALTER SYSTEM SET REMOTE_DEPENDENCIES_MODE = {SIGNATURE | TIMESTAMP}
```

If the REMOTE_DEPENDENCIES_MODE parameter is not specified, either in the init.ora parameter file or using the ALTER SESSION or ALTER SYSTEM statements, TIMESTAMP is the default value. Therefore, unless you explicitly use the REMOTE_DEPENDENCIES_MODE parameter, or the appropriate DDL statement, your server is operating using the time-stamp dependency mode.

When you use REMOTE_DEPENDENCIES_MODE=SIGNATURE:

- If you change the initial value of a parameter of a remote procedure, then the local procedure calling the remote procedure is not invalidated. If the call to the remote procedure does not supply the parameter, then the initial value is used. In this case, because invalidation and recompilation does not automatically occur, the old initial value is used. To see the new initial values, recompile the calling procedure manually.
- If you add an overloaded procedure in a package (a procedure with the same name as an existing one), then local procedures that call the remote procedure are not invalidated. If it turns out that this overloading results in a rebinding of existing calls from the local procedure under the time-stamp mode, then this rebinding does not happen under the RPC signature mode, because the local procedure does not get invalidated. You must recompile the local procedure manually to achieve the rebinding.
- If the types of parameters of an existing package procedure are changed so that the new types have the same shape as the old ones, then the local calling procedure is not invalidated or recompiled automatically. You must recompile the calling procedure manually to get the semantics of the new type.

Topics:

- Dependency Resolution
- Suggestions for Managing Dependencies

Dependency Resolution

When REMOTE_DEPENDENCIES_MODE = TIMESTAMP (the default value), dependencies among program units are handled by comparing time stamps at run

time. If the time stamp of a called remote procedure does not match the time stamp of the called procedure, then the calling (dependent) unit is invalidated and must be recompiled. In this case, if there is no local PL/SQL compiler, then the calling application cannot proceed.

In the time-stamp dependency mode, RPC signatures are not compared. If there is a local PL/SQL compiler, then recompilation happens automatically when the calling procedure is run.

When `REMOTE_DEPENDENCIES_MODE = SIGNATURE`, the recorded time stamp in the calling unit is first compared to the current time stamp in the called remote unit. If they match, then the call proceeds. If the time stamps do not match, then the RPC signature of the called remote subprogram, as recorded in the calling subprogram, is compared with the current RPC signature of the called subprogram. If they do not match (using the criteria described in the section "Changing Data Type Classes of Parameters" on page 18-16), then an error is returned to the calling session.

Suggestions for Managing Dependencies

Follow these guidelines for setting the `REMOTE_DEPENDENCIES_MODE` parameter:

- Server-side PL/SQL users can set the parameter to `TIMESTAMP` (or let it default to that) to get the time-stamp dependency mode.
- Server-side PL/SQL users can use RPC-signature dependency mode if they have a distributed system and they want to avoid possible unnecessary recompilations.
- Client-side PL/SQL users must set the parameter to `SIGNATURE`. This allows:
  - Installation of applications at client sites without recompiling procedures.
  - Ability to upgrade the server, without encountering time stamp mismatches.
- When using RPC signature mode on the server side, add procedures to the end of the procedure (or function) declarations in a package specification. Adding a procedure in the middle of the list of declarations can cause unnecessary invalidation and recompilation of dependent procedures.

Shared SQL Dependency Management

In addition to managing dependencies among schema objects, Oracle Database also manages dependencies of each shared SQL area in the shared pool. If a table, view, synonym, or sequence is created, altered, or dropped, or a procedure or package specification is recompiled, all dependent shared SQL areas are invalidated. At a subsequent execution of the cursor that corresponds to an invalidated shared SQL area, Oracle Database reparses the SQL statement to regenerate the shared SQL area.

19

Edition-Based Redefinition

Edition-based redefinition enables you to upgrade the database component of an application while it is in use, thereby minimizing or eliminating down time.

To upgrade an application while it is in use, you copy the database objects that comprise the application and redefine the copied objects in isolation. Your changes do not affect users of the application—they continue to run the unchanged application. When you are sure that your changes are correct, you make the upgraded application available to all users.

Using edition-based redefinition means using one or more of its component features. The features you use, and the down time, depend on these factors:

- What kind of database objects you redefine
- How available the database objects must be to users while you are redefining them
- Whether you make the upgraded application available to some users while others continue to use the older version of the application

You always use the **edition** feature to copy the database objects and redefine the copied objects in isolation; that is why the procedure that this chapter describes for upgrading applications online is called edition-based redefinition.

If the object type of every object you will redefine is **editionable** (defined in "Editionable and Noneditionable Schema Object Types" on page 19-3), the edition is the only feature you use.

Table is not an editionable type. If you change the structure of one or more tables, you also use the **editioning view** feature.

If other users must be able to change data in the tables while you are changing their structure, you also use **forward crossedition triggers**. If the pre- and post-upgrade applications will be in ordinary use at the same time (**hot rollover**), you also use **reverse crossedition triggers**. Crossedition triggers are not a permanent part of the application—you drop them when all users are using the post-upgrade application.

Topics:

- Editions
- Editioning Views
- Crossedition Triggers
- Displaying Information About Editions, Editioning Views, and Crossedition Triggers
- Using Edition-Based Redefinition to Upgrade an Application

Editions

Editions are nonschema objects; as such, they do not have owners. Editions are created in a single namespace, and multiple editions can coexist in the database.

The database must have at least one edition. Every newly created or upgraded Oracle Database starts with one edition named `ora$base`.

Topics:

- Editioned and Noneditioned Objects
- Creating an Edition
- Inherited and Actual Objects
- Making an Edition Available to Some Users
- Making an Edition Available to All Users
- Current Edition and Session Edition
- Retiring an Edition
- Dropping an Edition

Editioned and Noneditioned Objects

An **editioned object** is a schema object that has both an editionable type and an editions-enabled owner. (A schema object that has an editionable type but not an editions-enabled owner is **potentially editioned**.) An edition can have its own copy of an editioned object, in which case only the copy is visible to the edition.

A **noneditioned object** is a schema object that has a noneditionable type. An edition cannot have its own copy of a noneditioned object. A noneditioned object is identical in, and visible to, all editions.

An editioned object is uniquely identified by its `OBJECT_NAME`, `OWNER`, and `EDITION_NAME`. A noneditioned object is uniquely identified by its `OBJECT_NAME` and `OWNER`—its `EDITION_NAME` is `NULL`. (Strictly speaking, the `NAMESPACE` of an object is also required to uniquely identify the object, but you can ignore this fact, because any statement that references the object implicitly or explicitly specifies its `NAMESPACE`.)

You can display the `OBJECT_NAME`, `OWNER`, and `EDITION_NAME` of an object with the static data dictionary views `*_OBJECTS` and `*_OBJECTS_AE` (described in Table 19–1).

You need not know the `EDITION_NAME` of an object to refer to that object (and if you do know it, you cannot specify it). The context of the reference implicitly specifies the edition. If the context is a data definition language (DDL) statement, the edition is the current edition of the session that issued the command (for information about the current edition, see "Current Edition and Session Edition" on page 19-10). If the context is source code, the edition is the one in which the object is actual (see "Inherited and Actual Objects" on page 19-5).

Topics:

- Editionable and Noneditionable Schema Object Types
- Rules for Editioned Objects
- Enabling Editions for a User

Editionable and Noneditionable Schema Object Types

These schema objects types are **editionable**:

- `SYNONYM`
- `VIEW`
- All PL/SQL object types:
  - `FUNCTION`
  - `LIBRARY`
  - `PACKAGE` and `PACKAGE BODY`
  - `PROCEDURE`
  - `TRIGGER`
  - `TYPE` and `TYPE BODY`

All other schema object types are **noneditionable**. Table is an example of an noneditionable type.

A schema object of an editionable type is editioned if its owner is editions-enabled; otherwise, it is **potentially editioned**.

A schema object of a noneditionable type is always noneditioned, even if its owner is editions-enabled. A table is an example of an noneditioned object.

> **Note:** There is one exception to the rules: Although `SYNONYM` is an editionable type, a public synonym is a noneditioned object.

Rules for Editioned Objects

- A noneditioned object cannot depend on an editioned object.

  For example:

  - A public synonym cannot refer to an editioned object.
  - A function-based index cannot depend on an editioned function.
  - A materialized view cannot depend on an editioned view.
  - A table cannot have a column of a user-defined data type (collection or ADT) whose owner is editions-enabled.
  - A noneditioned subprogram cannot have a static reference to a subprogram whose owner is editions-enabled.

  For the reason for this rule, see "Actualizing Referenced Objects" on page 19-9.

- An ADT cannot be both editioned and evolved.

  For information about type evolution, see *Oracle Database Object-Relational Developer's Guide*.

- An editioned object cannot be the starting or ending point of a `FOREIGN KEY` constraint.

  The only editioned object that this rule affects is an editioned view. An editioned view can be either an ordinary view or an editioning view.

Enabling Editions for a User

To enable editions for a user, use the `ENABLE EDITIONS` clause of either the `CREATE USER` or `ALTER USER` statement.

The `EDITIONS_ENABLED` column of the static data dictionary view `*_USERS` shows which users have editions enabled.

Enabling editions is retroactive and irreversible. When a user is editions-enabled, every editionable-type object that the user has owned or will own is an editioned object. You cannot enable editions for a user who owns a potentially editioned object with noneditioned dependents unless you specify `FORCE`:

```
ALTER USER user_name ENABLE EDITIONS FORCE;
```

The preceding statement enables editions for the specified user and invalidates noneditioned dependents of editioned objects.

`FORCE` is useful in the following situation: You must editions-enable users `A` and `B`. User `A` owns potentially editioned objects `a1` and `a2`. User `B` owns potentially editioned objects `b1` and `b2`. Object `a1` depends on object `b1`. Object `b2` depends on object `a2`. Editions-enable users `A` and `B` like this:

1. Using `FORCE`, enable editions for user `A`:

   ```
   ALTER USER A ENABLE EDITIONS FORCE;
   ```

   Now `a1` and `a2` are editioned objects, and `b2` (which depends on `a2`) is invalid.

2. Enable editions for user `B`:

   ```
   ALTER USER B ENABLE EDITIONS;
   ```

3. Recompile `b2`, using the appropriate `ALTER` statement with `COMPILE`. For a PL/SQL object, also specify `REUSE SETTINGS`.

   For example, if `b2` is a procedure, use this statement:

   ```
   ALTER PROCEDURE b2 COMPILE REUSE SETTINGS
   ```

   For information about the `ALTER` statements for PL/SQL objects, see *Oracle Database PL/SQL Language Reference*.

   For information about the `ALTER` statements for SQL objects, see *Oracle Database SQL Language Reference*.

`FORCE` is unnecessary in the following situation: You must editions-enable user `C`, who owns potentially editioned object `c1`. Object `c1` has dependent `d1`, a potentially editioned object owned by user `D`. User `D` owns no potentially editioned objects that have dependents owned by `C`. If you editions-enable `D` first, making `d1` an editioned object, then you can editions-enable `C` without violating the rule that a noneditioned object cannot depend on an editioned object.

You cannot enable editions for a user who owns one or more evolved ADTs. Trying to do so causes error ORA-38820. If an ADT has no table dependents, you can use the `ALTER TYPE RESET` statement to reset its version to 1, so that it is no longer considered to be evolved. (Resetting the version of an ADT to 1 invalidates its dependents.)

See Also:

- *Oracle Database SQL Language Reference* for the syntax of the `CREATE USER`and `ALTER USER` statements
- *Oracle Database Reference* for more information about the static data dictionary views `*_USERS`

Creating an Edition

To create an edition, use the SQL statement `CREATE EDITION`.

You must create the edition as the child of an existing edition. The parent of the first edition created with a `CREATE EDITION` statement is `ora$base`. This statement creates the edition `e2` as the child of `ora$base`:

```
CREATE EDITION e2
```

(Example 19–1 and others use the preceding statement.)

At Release 11.2, an edition can have at most one child.

The **descendents** of an edition are its child, its child's child, and so on. The **ancestors** of an edition are its parent, its parent's parent, and so on. The **root edition** has no parent, and a **leaf edition** has no child.

> **See Also:** *Oracle Database SQL Language Reference* for information about the `CREATE EDITION` statement, including the privileges required to use it

Inherited and Actual Objects

Each database session uses exactly one edition at a time. Upon creation, a child edition inherits from its parent edition all editioned objects in the database that are visible in the parent edition. Each **inherited object** is visible in the child edition.

An inherited object is **copied on change** or **actualized**; that is, when a user of the child edition references an inherited object in a DDL statement (other than `DROP`), the inherited object is copied and the DDL statement affects only the copy—the **actual object**. The unchanged object in the parent edition is no longer visible in the child edition.

> **Note:** When the DDL statement `CREATE OR REPLACE` *`object`* has no effect, it does not actualize *`object`* (for details, see "Invalidation of Dependent Objects" on page 18-5). The DDL statement `ALTER` *`object`* `COMPILE` always actualizes *`object`*.

Example 19–1 creates a procedure named `hello` in the edition `ora$base`, and then creates the edition `e2` as a child of `ora$base`. When `e2` invokes `hello`, it invokes the inherited procedure. Then `e2` changes `hello`, actualizing it. The procedure `hello` in the edition `ora$base` remains unchanged, and is no longer visible in `e2`. Now when `e2` invokes `hello`, it invokes the actual procedure.

Example 19–1 Inherited and Actual Objects

1. Create procedure in parent edition:

```
CREATE OR REPLACE PROCEDURE hello IS
  BEGIN
```

```
    DBMS_OUTPUT.PUT_LINE('Hello, edition 1.');
  END hello;
/
```

2. Invoke procedure in parent edition:

```
BEGIN hello(); END;
/
```

Result:

```
Hello, edition 1.

PL/SQL procedure successfully completed.
```

3. Create child edition:

```
CREATE EDITION e2;
```

4. Use child edition:

```
ALTER SESSION SET EDITION = e2;
```

For information about `ALTER SESSION SET EDITION`, see "Changing Your Session Edition" on page 19-11.

5. In child edition, invoke procedure:

```
BEGIN hello(); END;
/
```

Child edition inherits procedure from parent edition. Child edition invokes inherited procedure. Result:

```
Hello, edition 1.

PL/SQL procedure successfully completed.
```

6. Change procedure in child edition:

```
CREATE OR REPLACE PROCEDURE hello IS
  BEGIN
    DBMS_OUTPUT.PUT_LINE('Hello, edition 2.');
  END hello;
/
```

Child changes only its own copy of procedure. Child's copy is an actual object.

7. Invoke procedure:

```
BEGIN hello(); END;
/
```

Child invokes its own copy, the actual procedure:

```
Hello, edition 2.

PL/SQL procedure successfully completed.
```

8. Return to parent:

```
ALTER SESSION SET EDITION = ora$base;
```

9. Invoke procedure and see that it has not changed:

```
BEGIN hello(); END;
/
```

Result:

```
Hello, edition 1.

PL/SQL procedure successfully completed.
```

Topics:

- Dropping Inherited Objects
- Actualizing Referenced Objects

Dropping Inherited Objects

If a user of the child edition drops an inherited object, that object is no longer visible in the child edition, but it continues to be visible in the parent edition.

Example 19–2 creates a procedure named `goodbye` in the edition `ora$base`, and then creates edition `e2` as a child of `ora$base`. After `e2` drops `goodbye`, it can no longer invoke it, but `ora$base` can still invoke it. (For more information about the `DROP PROCEDURE` statement, including the privileges required to use it, see *Oracle Database PL/SQL Language Reference*.)

Example 19–2 Dropping an Inherited Object

1. Create procedure in edition `ora$base`:

```
CREATE OR REPLACE PROCEDURE goodbye IS
  BEGIN
    DBMS_OUTPUT.PUT_LINE('Good-bye!');
  END goodbye;
/
```

2. Invoke procedure:

```
BEGIN goodbye; END;
/
```

Result:

```
Good-bye!

PL/SQL procedure successfully completed.
```

3. Create edition `e2` as a child of `ora$base`:

```
CREATE EDITION e2;
```

4. Use edition `e2`:

```
ALTER SESSION SET EDITION = e2;
```

`ALTER SESSION SET EDITION` must be a top-level SQL statement. For more information, see "Changing Your Session Edition" on page 19-11.

5. In `e2`, invoke procedure:

```
BEGIN goodbye; END;
/
```

`e2` invokes inherited procedure:

```
Good-bye!

PL/SQL procedure successfully completed.
```

6. In e2, drop procedure:

```
DROP PROCEDURE goodbye;
```

7. In e2, try to invoke dropped procedure:

```
BEGIN goodbye; END;
/
```

Result:

```
BEGIN goodbye; END;
      *
ERROR at line 1:
ORA-06550: line 1, column 7:
PLS-00201: identifier 'GOODBYE' must be declared
ORA-06550: line 1, column 7:
PL/SQL: Statement ignored
```

8. Return to parent:

```
ALTER SESSION SET EDITION = ora$base;
```

9. In parent, invoke procedure:

```
BEGIN goodbye; END;
/
```

Result:

```
Good-bye!

PL/SQL procedure successfully completed.
```

Because e2 dropped the procedure goodbye:

- Its descendents do not inherit the procedure goodbye.
- No object named goodbye is visible in e2, so e2 can create an object named goodbye of any editionable type. If e2 creates this object, the descendents of e2 inherit that object.

In Example 19–3, e2 creates a function named goodbye and then an edition named e3 as a child of e2. When e3 tries to invoke the *procedure* goodbye (which e2 dropped), an error occurs, but e3 successfully invokes the *function* goodbye (which e2 created).

Example 19–3 Creating an Object with the Name of a Dropped Inherited Object

1. Return to e2:

```
ALTER SESSION SET EDITION = e2;
```

For information about ALTER SESSION SET EDITION, see "Changing Your Session Edition" on page 19-11.

2. In e2, create function named goodbye:

```
CREATE OR REPLACE FUNCTION goodbye
  RETURN BOOLEAN
IS
```

```
BEGIN
  RETURN(TRUE);
END goodbye;
/
```

3. Create edition `e3`:

```
CREATE EDITION e3 AS CHILD OF e2;
```

4. Use edition `e3`:

```
ALTER SESSION SET EDITION = e3;
```

5. In `e3`, try to invoke *procedure* `goodbye`:

```
BEGIN
  goodbye;
END;
/
```

Result:

```
  goodbye;
  *
ERROR at line 2:
ORA-06550: line 2, column 3:
PLS-00221: 'GOODBYE' is not a procedure or is undefined
ORA-06550: line 2, column 3:
PL/SQL: Statement ignored
```

6. In `e3`, invoke *function* `goodbye`:

```
BEGIN
  IF goodbye THEN
    DBMS_OUTPUT.PUT_LINE('Good-bye!');
  END IF;
END;
/
```

Result:

```
Good-bye!

PL/SQL procedure successfully completed.
```

Actualizing Referenced Objects

When a referenced object is actualized in an edition, all of its dependents (direct and indirect) that are not yet actualized in that edition become actualized in that edition in an invalid state. Therefore, an editioned object cannot have dependents that cannot be actualized. In other words, a noneditioned object cannot depend on an editioned object (for examples, see "Rules for Editioned Objects" on page 19-3).

When an invalid object is referenced, the database automatically validates it, which requires name resolution. The database looks for the object name first in the current edition, then in the parent edition, and so on.

> **See Also:** Chapter 18, "Schema Object Dependency," for general information about dependencies among schema objects, including invalidation, revalidation, and name resolution

Making an Edition Available to Some Users

As the creator of the edition, you automatically have the `USE` privilege `WITH GRANT OPTION` on it. To grant the `USE` privilege on the edition to other users, use the SQL statement `GRANT USE ON EDITION`. For information about the `GRANT` statement, see *Oracle Database SQL Language Reference*.

Making an Edition Available to All Users

To make an edition available to all users, either:

- Grant the `USE` privilege on the edition to `PUBLIC`:

  ```
  GRANT USE ON EDITION edition_name TO PUBLIC
  ```

  For information about the `GRANT` statement, see *Oracle Database SQL Language Reference*.

- Make the edition the database default edition:

  ```
  ALTER DATABASE DEFAULT EDITION = edition_name
  ```

  This has the side effect of granting the `USE` privilege on *`edition_name`* to `PUBLIC`.

  For information about the `ALTER DATABASE` statement, see *Oracle Database SQL Language Reference*.

Current Edition and Session Edition

Each database session uses exactly one edition at a time. The edition that a database session is using at any one time is called its **current edition**. When a database session begins, its current edition is its **session edition**, which is the edition in which it begins. If you change the session edition, the current edition changes to the same thing. However, there are situations in which the current edition and session edition differ.

Topics:

- Your Initial Session Edition
- Changing Your Session Edition
- Displaying the Names of the Current and Session Editions
- When the Current Edition Might Differ from the Session Edition

Your Initial Session Edition

When you connect to the database, you can specify your initial session edition. Your initial session edition can be any edition on which you have the `USE` privilege. To see the names of the editions that are available to you, use this query:

```
SELECT EDITION_NAME FROM ALL_EDITIONS;
```

How you specify your initial session edition at connection time depends on how you connect to the database—see the documentation for your interface.

See Also:

- *Oracle Database Administrator's Guide* for information about setting the database default edition
- *SQL*Plus User's Guide and Reference* for information about connecting to the database with SQL*Plus
- *Oracle Call Interface Programmer's Guide* for information about connecting to the database with Oracle Call Interface (OCI)
- *Oracle Database JDBC Developer's Guide* for information about connecting to the database with JDBC

As of Oracle Database 11*g* Release 2 (11.2.0.2), if you do not specify your session edition at connection time, then:

- If you use a database service to connect to the database, and an initial session edition was specified for that service, then the initial session edition for the service is your initial session edition.
- Otherwise, your initial session edition is the database default edition.

As of Release 11.2.0.2, when you create or modify a database service, you can specify its initial session edition.

To create or modify a database service, Oracle recommends using the `srvctl add service` or `srvctl modify service` command. To specify the default initial session edition of the service, use the `-t` option.

Alternatively, you can create or modify a database service with the `DBMS_SERVICE.CREATE_SERVICE` or `DBMS_SERVICE.MODIFY_SERVICE` procedure, and specify the default initial session edition of the service with the `EDITION` attribute.

Note: As of Oracle Database 11*g* Release 2 (11.2.0.1), the `DBMS_SERVICE.CREATE_SERVICE` and `DBMS_SERVICE.MODIFY_SERVICE` procedures are deprecated in databases managed by Oracle Clusterware and Oracle Restart.

See Also:

- *Oracle Database Administrator's Guide* for information about the `-t` option of the `srvctl add service` command
- *Oracle Database Administrator's Guide* for information about the `-t` option of the `srvctl modify service` command
- *Oracle Database PL/SQL Packages and Types Reference* for information about the `EDITION` attribute of the `DBMS_SERVICE.CREATE_SERVICE` procedure
- *Oracle Database PL/SQL Packages and Types Reference* for information about the `EDITION` attribute of the `DBMS_SERVICE.MODIFY_SERVICE` procedure

Changing Your Session Edition

After connecting to the database, you can change your session edition with the SQL statement `ALTER SESSION SET EDITION`. You can change your session edition to any edition on which you have the `USE` privilege. When you change your session edition, your current edition changes to the same thing.

These statements from Example 19–1 and Example 19–2 change the session edition (and current edition) first to e2 and later to ora$base:

```
ALTER SESSION SET EDITION = e2
...
ALTER SESSION SET EDITION = ora$base
```

Note: ALTER SESSION SET EDITION must be a top-level SQL statement. To defer an edition change (in a logon trigger, for example), use the DBMS_SESSION.SET_EDITION_DEFERRED procedure.

See Also:

- *Oracle Database SQL Language Reference* for more information about the ALTER SESSION SET EDITION statement
- *Oracle Database PL/SQL Packages and Types Reference* for more information about the DBMS_SESSION.SET_EDITION_DEFERRED procedure

Displaying the Names of the Current and Session Editions

This statement returns the name of the current edition:

```
SELECT SYS_CONTEXT('USERENV', 'CURRENT_EDITION_NAME') FROM DUAL;
```

This statement returns the name of the session edition:

```
SELECT SYS_CONTEXT('USERENV', 'SESSION_EDITION_NAME') FROM DUAL;
```

See Also: *Oracle Database SQL Language Reference* for more information about the SYS_CONTEXT function

When the Current Edition Might Differ from the Session Edition

The current edition might differ from the session edition in these situations:

- A crossedition trigger fires.

  For details, see "Crossedition Trigger Interaction with Editions" on page 19-19.

- You run a statement by calling the DBMS_SQL.PARSE procedure, specifying the edition in which the statement is to run, as in Example 19–4.

  While the statement is running, the current edition is the specified edition, but the session edition does not change. For information about the DBMS_SQL.PARSE procedure, see *Oracle Database PL/SQL Packages and Types Reference*.

Example 19–4 creates a function that returns the names of the session edition and current edition. Then it creates a child edition, which invokes the function twice. The first time, the session edition and current edition are the same. The second time, they are not, because a different edition is passed as a parameter to the DBMS_SQL.PARSE procedure.

Example 19–4 Current Edition Differs from Session Edition

1. Create function that returns the names of the session edition and current edition:

```
CREATE OR REPLACE FUNCTION session_and_current_editions
  RETURN VARCHAR2
IS
```

```
BEGIN
  RETURN
  'Session: '|| SYS_CONTEXT('USERENV', 'SESSION_EDITION_NAME') ||
  ' / ' ||
  'Current: '|| SYS_CONTEXT('USERENV', 'CURRENT_EDITION_NAME');
END session_and_current_editions;
/
```

2. Create child edition:

```
CREATE EDITION e2 AS CHILD OF ora$base;
```

3. Use child edition:

```
ALTER SESSION SET EDITION = e2;
```

4. Invoke function:

```
BEGIN
  DBMS_OUTPUT.PUT_LINE (session_and_current_editions());
END;
/
```

Result:

```
Session: E2 / Current: E2

PL/SQL procedure successfully completed.
```

5. Invoke function again:

```
DECLARE
  c     NUMBER := DBMS_SQL.OPEN_CURSOR();
  v     VARCHAR2(200);
  dummy NUMBER;
  stmt  CONSTANT VARCHAR2(32767)
    := 'SELECT session_and_current_editions() FROM DUAL';
BEGIN
  DBMS_SQL.PARSE (c => c,
                  statement => stmt,
                  language_flag => DBMS_SQL.NATIVE,
                  edition => 'ora$base');

  DBMS_SQL.DEFINE_COLUMN (c, 1, v, 200);
  dummy := DBMS_SQL.EXECUTE_AND_FETCH (c, true);
  DBMS_SQL.COLUMN_VALUE (c, 1, v);
  DBMS_SQL.CLOSE_CURSOR(c);
  DBMS_OUTPUT.PUT_LINE (v);
END;
/
```

Result:

```
Session: E2 / Current: ORA$BASE

PL/SQL procedure successfully completed.
```

Retiring an Edition

After making a new edition (an upgraded application) available to all users, you want to retire the old edition (the original application), so that no user except `SYS` can use the retired edition.

> **Note:** If the old edition is the database default edition, make another edition the database default edition before you retire the old edition:
>
> ```
> ALTER DATABASE DEFAULT EDITION = edition_name
> ```
>
> For information about the `ALTER DATABASE` statement, see *Oracle Database SQL Language Reference*.

To retire an edition, you must revoke the `USE` privilege on the edition from every grantee. To list the grantees, use this query, where `:e` is a placeholder for the name of the edition to be dropped:

```
SELECT GRANTEE, PRIVILEGE
FROM DBA_TAB_PRIVS
WHERE TABLE_NAME = :e
/
```

For information about the `REVOKE` statement, see *Oracle Database SQL Language Reference*.

Dropping an Edition

> **Note:** If the edition includes crossedition triggers, see "Dropping the Crossedition Triggers" on page 19-26 before you drop the edition.

To drop an edition, use the `DROP EDITION` statement, described in *Oracle Database SQL Language Reference*. If the edition has actual objects, you must specify the `CASCADE` clause, which drops the actual objects.

If a `DROP EDITION` *`edition`* `CASCADE` statement is interrupted before finishing normally (from a power failure, for example), the static data dictionary view `*_EDITIONS` shows that the value of `USABLE` for *`edition`* is `NO`. The only operation that you can perform on such an unusable *`edition`* is `DROP EDITION CASCADE`.

You drop an edition in these situations:

- You want to roll back the application upgrade.
- (Optional) You have retired the edition.

You can drop an edition only if all of these statements are true:

- The edition is either the root edition or a leaf edition.
- If the edition is the root, it has no objects that its descendents inherit. (That is, each object inherited from the root edition was either actualized or dropped.)
- The edition is not in use (that is, it is not the current edition or session edition of a session).
- The edition is not the database default edition.

To explicitly actualize an inherited object in the child edition:

1. Make the child edition your session edition.

   For instructions, see "Changing Your Session Edition" on page 19-11.

2. Recompile the object, using the appropriate `ALTER` statement with `COMPILE`. For a PL/SQL object, also specify `REUSE SETTINGS`.

   For example, this statement actualizes the procedure `p1`:

   ```
   ALTER PROCEDURE p1 COMPILE REUSE SETTINGS
   ```

   For information about the `ALTER` statements for PL/SQL objects, see *Oracle Database PL/SQL Language Reference*.

   For information about the `ALTER` statements for SQL objects, see *Oracle Database SQL Language Reference*.

   See Also:

   - *Oracle Database SQL Language Reference* for information about the `ALTER LIBRARY` statement
   - *Oracle Database SQL Language Reference* for information about the `ALTER VIEW` statement
   - *Oracle Database PL/SQL Language Reference* for information about the `ALTER FUNCTION` statement
   - *Oracle Database PL/SQL Language Reference* for information about the `ALTER PACKAGE` statement
   - *Oracle Database PL/SQL Language Reference* for information about the `ALTER PROCEDURE` statement
   - *Oracle Database PL/SQL Language Reference* for information about the `ALTER TRIGGER` statement
   - *Oracle Database PL/SQL Language Reference* for information about the `ALTER TYPE` statement

Editioning Views

The owner of an editioning view must be editions-enabled before the editioning view is created.

On a noneditioning view, the only type of trigger that you can define is an `INSTEAD OF` trigger. On an editioning view, you can define every type of trigger that you can define on a table (except crossedition triggers, which are temporary, and `INSTEAD OF` triggers). Therefore, and because they can be editioned, editioning views let you treat their base tables as if the base tables were editioned. However, you cannot add indexes or constraints to an editioning view; if your upgraded application requires new indexes or constraints, you must add them to the base table.

An editioning view selects a subset of the columns from a single base table and, optionally, provides aliases for them. In providing aliases, the editioning view maps physical column names (used by the base table) to logical column names (used by the application). An editioning view is like an API for a table.

There is no performance penalty for accessing a table through an editioning view, rather than directly. That is, if a SQL `SELECT`, `INSERT`, `UPDATE`, `DELETE`, or `MERGE` statement uses one or more editioning views, one or more times, and you replace each editioning view name with the name of its base table and adjust the column names if necessary, performance does not change.

The static data dictionary view `*_EDITIONING_VIEWS` describes every editioning view in the database that is visible (actual or inherited) in the session edition. `*_EDITIONING_VIEWS_AE` describes every actual object in every editioning view in the database, in every edition.

Topics:

- Creating an Editioning View
- Partition-Extended Editioning View Names
- Changing the 'Write-ability' of an Editioning View
- Replacing an Editioning View
- Dropping or Renaming the Base Table
- Adding Indexes and Constraints to the Base Table
- SQL Optimizer Index Hints

> **See Also:** *Oracle Database Reference* for more information about the static data dictionary views `*_EDITIONING_VIEWS` and `*_EDITIONING_VIEWS_AE`.

Creating an Editioning View

To create an editioning view, use the SQL statement `CREATE VIEW` with the keyword `EDITIONING`. To make the editioning view read-only, specify `WITH READ ONLY`; to make it read-write, omit `WITH READ ONLY`.

If an editioning view is **read-only**, users of the unchanged application can see the data in the base table, but cannot change it. The base table has **semi-availability**. Semi-availability is acceptable for applications such as online dictionaries, which users read but do not change. Make the editioning view read-only if you do not define crossedition triggers on the base table.

If an editioning view is **read-write**, users of the unchanged application can both see and change the data in the base table. The base table has **maximum availability**. Maximum availability is required for applications such as online stores, where users submit purchase orders. If you define crossedition triggers on the base table, make the editioning view read-write.

Because an editioning view must do no more than select a subset of the columns from the base table and provide aliases for them, the `CREATE VIEW` statement that creates an editioning view has restrictions. Violating the restrictions causes the creation of the view to fail, even if you specify `FORCE`.

> **See Also:** *Oracle Database SQL Language Reference* for more information about using the `CREATE VIEW` statement to create editioning views, including the restrictions

Partition-Extended Editioning View Names

An editioning view defined on a partitioned table can have a partition-extended name, with partition and subpartition names that refer to the partitions and subpartitions of the base table.

The data manipulation language (DML) statements that support partition-extended table names also support partition-extended editioning view names. These statements are:

- DELETE
- INSERT
- SELECT
- UPDATE

See Also: *Oracle Database SQL Language Reference* for information about referring to partitioned tables

Changing the 'Write-ability' of an Editioning View

To change an existing editioning view from read-only to read-write, use the SQL statement `ALTER VIEW READ WRITE`. To change an existing editioning view from read-write to read-only, use the SQL statement `ALTER VIEW READ ONLY`.

See Also: *Oracle Database SQL Language Reference* for more information about the `ALTER VIEW` statement

Replacing an Editioning View

To replace an editioning view, use the SQL statement `CREATE VIEW` with the `OR REPLACE` clause and the keyword `EDITIONING`.

You can replace an editioning view only with another editioning view. Any triggers defined on the replaced editioning view are retained.

Dropping or Renaming the Base Table

If you drop or rename the base table on which an editioning view is defined, the editioning view is not dropped, but the editioning view and its dependents become invalid. However, any triggers defined on the editioning view remain.

Adding Indexes and Constraints to the Base Table

If your upgraded application requires new indexes or constraints, you must add them to the base table. You cannot add them to the editioning view.

If the new indexes might negatively impact the old edition (the original application), make them invisible. In the crossedition triggers that must use the new indexes, specify them in `INDEX` hints.

When all users are using only the upgraded application:

- If the new indexes were used only by the crossedition triggers, drop them.
- If the new indexes are helpful in the upgraded application, make them visible.

See Also:

- "Guidelines for Application-Specific Indexes" on page 4-1 for information about when to use indexes
- *Oracle Database SQL Language Reference* for information about `INDEX` hints
- "SQL Optimizer Index Hints" on page 19-18

SQL Optimizer Index Hints

SQL optimizer index hints are specified in terms of the logical names of the columns participating in the index. Any SQL optimizer index hints specified on an editioning view using logical column names must be mapped to an index on the corresponding physical column in the base table.

> **See Also:** *Oracle Database SQL Language Reference* for information about using hints

Crossedition Triggers

The most important difference between crossedition triggers and noncrossedition triggers is how they interact with editions. A crossedition trigger is visible only in the edition in which it is actual, never in a descendent edition. Forward crossedition triggers move data from columns used by the old edition to columns used by the new edition; reverse crossedition triggers do the reverse.

Other important differences are:

- Crossedition triggers can be ordered with triggers defined on other tables, while noncrossedition triggers can be ordered only with other triggers defined on the same table.
- Crossedition triggers are temporary—you drop them after you have made the restructured tables available to all users.

Topics:

- Forward Crossedition Triggers
- Reverse Crossedition Triggers
- Crossedition Trigger Interaction with Editions
- Creating a Crossedition Trigger
- Transforming Data from Pre- to Post-Upgrade Representation
- Dropping the Crossedition Triggers

> **See Also:** *Oracle Database PL/SQL Language Reference* for general information about triggers

Forward Crossedition Triggers

The DML changes that you make to the table in the post-upgrade edition are written only to new columns or new tables, never to columns that users of pre-upgrade (ancestor) editions might be reading or writing. However, if the user of an ancestor edition changes the table data, the editioning view that you see must accurately reflect these changes. This is accomplished with forward crossedition triggers.

A forward crossedition trigger defines a **transform**, which is a rule for transforming an old row to one or more new rows. An **old row** is a row of data in the pre-upgrade representation. A **new row** is a row of data in the post-upgrade representation. The name of the trigger refers to the trigger itself and to the transform that the trigger defines.

Reverse Crossedition Triggers

If the pre- and post-upgrade editions will be in ordinary use at the same time (hot rollover), use reverse crossedition triggers to ensure that when users of the post-upgrade edition make changes to the table data, the changes are accurately reflected in the pre-upgrade editions.

Crossedition Trigger Interaction with Editions

The most important difference between crossedition triggers and noncrossedition triggers is how they interact with editions.

In this topic, the **current edition** is the edition in which the triggering DML statement runs. The current edition might differ from the session edition (for details, see "When the Current Edition Might Differ from the Session Edition" on page 19-12).

- Which Triggers Are Visible
- What Kind of Triggers Can Fire
- Firing Order
- Crossedition Trigger Execution

Which Triggers Are Visible

Editions inherit noncrossedition triggers in the same way that they inherit other editioned objects (see "Inherited and Actual Objects" on page 19-5).

Editions do not inherit crossedition triggers. A crossedition trigger might fire in response to a DML statement that another edition runs, but its name is visible only in the edition in which it was created. Therefore, an edition can reuse the name of a crossedition trigger created in an ancestor edition. Reusing the name of a crossedition trigger does not change the conditions under which the older trigger fires.

Crossedition triggers that appear in static data dictionary views are actual objects in the current edition.

What Kind of Triggers Can Fire

What kind of triggers can fire depends on the category of the triggering DML statement. The categories are:

- "Forward Crossedition Trigger SQL" on page 19-19
- "Reverse Crossedition Trigger SQL" on page 19-20
- "Application SQL" on page 19-20

> **Note:** The `APPEND` hint on a SQL `INSERT` statement does not prevent crossedition triggers from firing. For information about the `APPEND` hint, see *Oracle Database SQL Language Reference*.

Forward Crossedition Trigger SQL **Forward crossedition trigger SQL** is SQL that is executed in either of these ways:

- Directly from the body of a forward crossedition trigger

  This category includes SQL in an invoked subprogram only if the subprogram is local to the forward crossedition trigger.

- By invoking the `DBMS_SQL.PARSE` procedure with a non-`NULL` value for the `apply_crossedition_trigger` parameter

  The only valid non-`NULL` value for the `apply_crossedition_trigger` parameter is the unqualified name of a forward crossedition trigger. For more information about the `DBMS_SQL.PARSE` procedure, see *Oracle Database PL/SQL Packages and Types Reference*.

If a forward crossedition trigger invokes a subprogram in another compilation unit, the SQL in the subprogram is forward crossedition trigger SQL only if it is invoked by the `DBMS_SQL.PARSE` procedure with a non-`NULL` value for the `apply_crossedition_trigger` parameter.

Forward crossedition trigger SQL can fire only triggers that satisfy all of these conditions:

- They are forward crossedition triggers.
- They were created either in the current edition or in a descendent of the current edition.
- They explicitly follow the running forward crossedition trigger.

Reverse Crossedition Trigger SQL **Reverse crossedition trigger SQL** is SQL that is executed directly from the body of a reverse crossedition trigger. This category includes SQL in an invoked subprogram only if the subprogram is local to the reverse crossedition trigger.

Reverse crossedition trigger SQL can fire only triggers that satisfy all of these conditions:

- They are reverse crossedition triggers.
- They were created either in the current edition or in an ancestor of the current edition.
- They explicitly precede the running reverse crossedition trigger.

Application SQL **Application SQL** is all SQL except crossedition trigger SQL, including these DML statements:

- Dynamic SQL DML statements coded with the `DBMS_SQL` package (for information about these statements, see *Oracle Database PL/SQL Language Reference*).
- DML statements executed by Java stored procedures and external procedures (even when these procedures are invoked by `CALL` triggers)

Application SQL fires both noncrossedition and crossedition triggers, according to these rules:

Kind of Trigger	Conditions Under Which Trigger Can Fire
Noncrossedition	Trigger is both visible and enabled in the current edition.
Forward crossedition	Trigger was created in a descendent of the current edition.
Reverse crossedition	Trigger was created either in the current edition or in an ancestor of the current edition.

Firing Order

For a trigger to fire in response to a specific DML statement, the trigger must:

- Be the right kind (see "What Kind of Triggers Can Fire" on page 19-19)
- Satisfy the selection criteria (for example, the type of DML statement and the `WHEN` clause)
- Be enabled

For the triggers that meet these requirements, firing order depends on:

- FOLLOWS and PRECEDES Clauses
- Trigger Type and Edition

> **See Also:** *Oracle Database PL/SQL Language Reference* for general information about trigger firing order

FOLLOWS and PRECEDES Clauses When triggers A and B are to be fired at the same timing point, A fires before B fires if either of these is true:

- A explicitly precedes B.
- B explicitly follows A.

This rule is independent of conditions such as:

- Whether the triggers are enabled or disabled
- Whether the columns specified in the `UPDATE OF` clause are modified
- Whether the `WHEN` clauses are satisfied
- Whether the triggers are associated with the same kinds of DML statements (`INSERT`, `UPDATE`, or `DELETE`)
- Whether the triggers have overlapping timing points

The firing order of triggers that do not explicitly follow or precede each other is unpredictable.

Trigger Type and Edition For each timing point associated with a triggering DML statement, eligible triggers fire in this order. In categories 1 through 3, `FOLLOWS` relationships apply; in categories 4 and 5, `PRECEDES` relationships apply.

1. Noncrossedition triggers
2. Forward crossedition triggers created in the current edition
3. Forward crossedition triggers created in descendents of the current edition, in the order that the descendents were created (child, grandchild, and so on)
4. Reverse crossedition triggers created in the current edition
5. Reverse crossedition triggers created in the ancestors of the current edition, in the reverse order that the ancestors were created (parent, grandparent, and so on)

Crossedition Trigger Execution

A crossedition trigger runs using the edition in which it was created. Any code that the crossedition trigger calls (including package references, PL/SQL subprogram calls, and SQL statements) also runs in the edition in which the crossedition trigger was created.

If a PL/SQL package is actual in multiple editions, then the package variables and other state are private in each edition, even within a single session. Because each crossedition trigger and the code that it calls run using the edition in which the

crossedition trigger was created, the same session can instantiate two or more versions of the package, with the same name.

Creating a Crossedition Trigger

To create a crossedition trigger, you must be editions-enabled (for information about enabling editions for a user, see "Enabling Editions for a User" on page 19-4).

Create a crossedition trigger with the SQL statement `CREATE TRIGGER`, observing these rules:

- A crossedition trigger must be defined on a table, not a view.
- A crossedition trigger must be a DML trigger (simple or compound).

  The DML statement in a crossedition trigger body can be either a static SQL statement (described in *Oracle Database PL/SQL Language Reference*) or a native dynamic SQL statement (described in *Oracle Database PL/SQL Language Reference*).
- A crossedition trigger is forward unless you specify `REVERSE`. (Specifying `FORWARD` is optional.)
- The `FOLLOWS` clause is allowed only when creating a forward crossedition trigger or a noncrossedition trigger. (The `FOLLOWS` clause indicates that the trigger being created is to fire after the specified triggers fire.)
- The `PRECEDES` clause is allowed only when creating a reverse crossedition trigger. (The `PRECEDES` clause indicates that the trigger being created is to fire before the specified triggers fire.)
- The triggers specified in the `FOLLOWS` or `PRECEDES` clause must exist, but need not be enabled or successfully compiled.
- Like a noncrossedition trigger, a crossedition trigger is created in the enabled state unless you specify `DISABLE`. (Specifying `ENABLE` is optional.)

  Tip: Create crossedition triggers in the disabled state, and enable them after you are sure that they compile successfully. If you create them in the enabled state, and they fail to compile, the failure affects users of the existing application.

- The operation in a crossedition trigger body must be **idempotent** (that is, performing the operation multiple times is redundant; it does not change the result).

  See Also: *Oracle Database PL/SQL Language Reference* for more information about using the `CREATE TRIGGER` statement to create crossedition triggers

Coding the Forward Crossedition Trigger Body

The operation in the body of a forward crossedition trigger must be idempotent, because it is impossible to predict:

- The context in which the body will first run for an old row.

  The possibilities are:

  – When a user of an ancestor edition runs a DML statement that fires the trigger (a **serendipitous change**)

  – When you **apply the transform** that the trigger defines

For information about applying transforms, see "Transforming Data from Pre- to Post-Upgrade Representation" on page 19-24.

- How many times the body will run for each old row.

Topics:

- Handling Data Transformation Collisions
- Handling Changes to Other Tables

Handling Data Transformation Collisions If a forward crossedition trigger populates a new table (rather than new columns of a table), its body must handle data transformation collisions.

For example, suppose that a column of the new table has a `UNIQUE` constraint. A serendipitous change fires the forward crossedition trigger, which inserts a row in the new table. Later, another serendipitous change fires the forward crossedition trigger, or you apply the transform defined by the trigger. The trigger tries to insert a row in the new table, violating the `UNIQUE` constraint.

If your collision-handling strategy depends on why the trigger is running, you can determine the reason with the function `APPLYING_CROSSEDITION_TRIGGER`. When called directly from a trigger body, this function returns the `BOOLEAN` value `TRUE` if the trigger is running because of a serendipitous change and `FALSE` if the trigger is running because you are applying the transform. (`APPLYING_CROSSEDITION_TRIGGER` is defined in the package `DBMS_STANDARD`. It has no parameters.)

To ignore collisions and insert the rows that do not collide with existing rows, put the `IGNORE_ROW_ON_DUPKEY_INDEX` hint in the `INSERT` statement.

If you do not want to ignore such collisions, but want to know where they occur so that you can handle them, put the `CHANGE_DUPKEY_ERROR_INDEX` hint in the `INSERT` or `UPDATE` statement, specifying either an index or set of columns. Then, when a unique key violation occurs for that index or set of columns, ORA-38911 is reported instead of ORA-00001. You can write an exception handler for ORA-38911.

Note: Although they have the syntax of hints, `IGNORE_ROW_ON_DUPKEY_INDEX` and `CHANGE_DUPKEY_ERROR_INDEX` are mandates. The optimizer always uses them.

Example 19–5 creates a crossedition trigger that uses the `APPLYING_CROSSEDITION_TRIGGER` function and the `IGNORE_ROW_ON_DUPKEY_INDEX` and `CHANGE_DUPKEY_ERROR_INDEX` hints to handle data transformation collisions. The trigger transforms old rows in `table1` to new rows in `table2`. The tables were created as follows:

```
CREATE TABLE table1 (key NUMBER, value VARCHAR2(20));

CREATE TABLE table2 (key NUMBER, value VARCHAR2(20), last_updated TIMESTAMP);
CREATE UNIQUE INDEX i2 on table2(key);
```

Example 19–5 Crossedition Trigger that Handles Data Transformation Collisions

```
CREATE OR REPLACE TRIGGER trigger1
  BEFORE INSERT OR UPDATE ON table1
  FOR EACH ROW
  CROSSEDITION
DECLARE
  row_already_present  EXCEPTION;
```

```
  PRAGMA EXCEPTION_INIT(row_already_present, -38911);
BEGIN
  IF APPLYING_CROSSEDITION_TRIGGER THEN
    /* Trigger is running because of serendipitous change.
       Insert new row into table2 unless it is already there. */
    INSERT /*+ IGNORE_ROW_ON_DUPKEY_INDEX(table2(key)) */
    INTO table2
    VALUES(:new.key, :new.value, to_date('1900-01-01', 'YYYY-MM-DD'));
  ELSE
    /* Trigger is running because you are applying transform.
       If tranform has not yet inserted new row in table2, insert new row;
       otherwise, update new row. */
    BEGIN
      INSERT /*+ CHANGE_DUPKEY_ERROR_INDEX(table2(key)) */
      INTO table2
      VALUES(:new.key, :new.value, SYSTIMESTAMP);
    EXCEPTION WHEN row_already_present THEN
      UPDATE table2
      SET value = :new.value, last_updated = SYSTIMESTAMP
      WHERE key = :new.key;
    END;
  END IF;
END;
/
```

See Also:

- *Oracle Database SQL Language Reference* for more information about `IGNORE_ROW_ON_DUPKEY_INDEX`
- *Oracle Database SQL Language Reference* for more information about `CHANGE_DUPKEY_ERROR_INDEX`
- *Oracle Database SQL Language Reference* for general information about hints

Handling Changes to Other Tables If the body of a forward crossedition trigger includes explicit SQL statements that change tables other than the one on which the trigger is defined, and if the rows of those tables do not have a one-to-one correspondence with the rows of the table on which the trigger is defined, then the body code must implement a locking mechanism that correctly handles these situations:

- Two or more users of ancestor editions simultaneously issue DML statements for the table on which the trigger is defined.
- At least one user of an ancestor edition issues a DML statement for the table on which the trigger is defined.

Transforming Data from Pre- to Post-Upgrade Representation

After redefining the database objects that comprise the application that you are upgrading (in the new edition), you must transform the application data from its pre-upgrade representation (in the old edition) to its post-upgrade representation (in the new edition). The rules for this transformation are called **transforms**, and they are defined by forward crossedition triggers. (For general information about forward crossedition triggers, see "Forward Crossedition Triggers" on page 19-18.)

Some old rows might have been transformed to new rows by **serendipitous changes**; that is, by changes that users of the pre-upgrade application made, which fired forward crossedition triggers. However, any rows that were not tranformed by

serendipitous changes are still in their pre-upgrade representation. To ensure that all old rows are transformed to new rows, you must **apply the transforms** that you defined on the tables that store the application data.

There are two ways to apply a transform:

- Fire the trigger that defines the transform on every row of the table, one row at a time.
- Instead of firing the trigger, run a SQL statement that does what the trigger would do, but faster, and then fire any triggers that follow that trigger.

  This second way is recommended if you have replaced an entire table or created a new table.

For either way of applying the transform, invoke either the `DBMS_SQL.PARSE` procedure or the subprograms in the `DBMS_PARALLEL_EXECUTE` package. The latter is recommended if you have a lot of data. The subprograms enable you to incrementally update the data in a large table in parallel, in two high-level steps:

1. Group sets of rows in the table into smaller chunks.
2. Apply the desired `UPDATE` statement to the chunks in parallel, committing each time you have finished processing a chunk.

The advantages are:

- You lock only one set of rows at a time, for a relatively short time, instead of locking the entire table.
- You do not lose work that has been done if something fails before the entire operation finishes.

For both the `DBMS_SQL.PARSE` procedure and the `DBMS_PARALLEL_EXECUTE` subprograms, the actual parameter values for `apply_crossedition_trigger`, `fire_apply_trigger`, and `sql_stmt` are the same:

- For `apply_crossedition_trigger`, specify the name of the forward crossedition trigger that defines the transform to be applied.
- To fire the trigger on every row of the table, one row at a time:
  - For the value of `fire_apply_trigger`, specify `TRUE`.
  - For `sql_stmt`, supply a SQL statement whose only significant effect is to select the forward crossedition trigger to be fired; for example, an `UPDATE` statement that sets some column to its own existing value in each row.
- To run a SQL statement that does what the trigger would do, and then fire any triggers that follow that trigger:
  - For the value of `fire_apply_trigger`, specify `FALSE`.
  - For `sql_stmt`, supply a SQL statement that does what the forward crossedition trigger would do, but faster—for example, a PL/SQL anonymous block that calls one or more PL/SQL subprograms.

    See Also:

    - *Oracle Database PL/SQL Packages and Types Reference* for more information about the `DBMS_SQL.PARSE` procedure
    - *Oracle Database PL/SQL Packages and Types Reference* for more information about the `DBMS_PARALLEL_EXECUTE` package

Preventing Lost Updates

To prevent lost updates when applying a transform, use this procedure:

1. Enable crossedition triggers.
2. Wait until pending changes to the affected tables are either committed or rolled back.

   Use the procedure `DBMS_UTILITY.WAIT_ON_PENDING_DML`, described in *Oracle Database PL/SQL Packages and Types Reference*.
3. Apply the transform.

Note: This scenario, where the forward crossedition trigger changes only the table on which it is defined, is sufficient to illustrate the risk. Suppose that Session One issues an `UPDATE` statement against the table when the crossedition trigger is not yet enabled; and that Session Two then enables the crossedition trigger and immediately applies the transformation.

A race condition can now occur when both Session One and Session Two will change the same row (row *n*). Chance determines which session reaches row *n* first. Both updates succeed, even if the session that reaches row *n* second must wait until the session that reached it first commits its change and releases its lock.

The problem occurs when Session Two wins the race. Because its SQL statement was compiled after the trigger was enabled, the program that implements the statement also implements the trigger action; therefore, the intended post-upgrade column values are set for row *n*. Now Session One reaches row *n*, and because its SQL statement was compiled before the trigger was enabled, the program that implements the statement does not implement the trigger action. Therefore, the values that Session Two set in the post-upgrade columns do not change—they reflect the values that the source columns had before Session One updated row *n*. That is, the intended side-effect of Session One's update is lost.

Dropping the Crossedition Triggers

To drop a crossedition trigger, use the `DROP TRIGGER` statement, described in *Oracle Database PL/SQL Language Reference*. Alternatively, you can drop crossedition triggers by dropping the edition in which they are actual, by using the `DROP EDITION` statement with the `CASCADE` clause. For information about dropping editions, see "Dropping an Edition" on page 19-14.

You drop crossedition triggers in these situations:

- You are rolling back the application upgrade (dropping the post-upgrade edition).

  Before dropping the post-upgrade edition, you must disable or drop any constraints on the new columns.
- You have finished the application upgrade and made the post-upgrade edition available to all users.

  When all sessions are using the post-upgrade edition, you can drop the forward crossedition triggers. However, before dropping the reverse crossedition triggers, you must disable or drop any constraints on the old columns.

To disable or drop constraints, use the ALTER TABLE statement with the DISABLE CONSTRAINT or DROP CONSTRAINT clause. For information about the ALTER TABLE statement, see *Oracle Database SQL Language Reference*.

Displaying Information About Editions, Editioning Views, and Crossedition Triggers

Table 19–1 and Table 19–2 describe the static data dictionary views that display information about editions and editioning views, respectively. For more information about a view in either table, see *Oracle Database Reference*.

The static data dictionary views that display information about triggers are described in *Oracle Database Reference*. Crossedition triggers that appear in static data dictionary views are actual objects in the session edition.

Child cursors cannot be shared if the set of crossedition triggers that might run differs. The dynamic performance views V$SQL_SHARED_CURSOR and GV$SQL_SHARED_CURSOR have a CROSSEDITION_TRIGGER_MISMATCH column that tells whether this is true. For information about V$SQL_SHARED_CURSOR, see *Oracle Database Reference*.

Table 19–1 *_ Dictionary Views with Edition Information

View	Description
*_EDITIONS	Describe every edition in the database.
*_EDITION_COMMENTS	Show the comments associated with every edition in the database.
*_OBJECTS	Describe every object in the database that is visible (actual or inherited) in the session edition.
*_OBJECTS_AE	Describe every actual object in the database, in every edition.
*_ERRORS	Describe every error in the database in the session edition.
*_ERRORS_AE	Describe every error in the database, in every edition.
*_USERS	Describe every user in the database. Useful for showing which users have editions enabled.
*_SERVICES	Describe every service in the database. As of Oracle Database 11*g* Release 2 (11.2.0.2), the EDITIONS column shows the default initial session edition.

Note: *_OBJECTS and *_OBJECTS_AE include dependent objects that are invalidated by operations in Table 18–2 only after one of the following:

- A reference to the object (either during compilation or execution)
- An invocation of DBMS_UTILITY.COMPILE_SCHEMA (described in *Oracle Database PL/SQL Packages and Types Reference*)
- An invocation of any UTL_RECOMP subprogram (described in *Oracle Database PL/SQL Packages and Types Reference*)

*Table 19–2 *_ Dictionary Views with Editioning View Information*

View	Description
*_VIEWS	Describes every view in the database that is visible (actual or inherited) in the session edition, including editioning views.
*_EDITIONING_VIEWS	Describes every editioning view in the database that is visible (actual or inherited) in the session edition. Useful for showing relationships between editioning views and their base tables. Join with *_OBJECTS_AE for additional information.
*_EDITIONING_VIEWS_AE	Describes every actual object in every editioning view in the database, in every edition.
*_EDITIONING_VIEW_COLS	Describes the columns of every editioning view in the database that is visible (actual or inherited) in the session edition. Useful for showing relationships between the columns of editioning views and the table columns to which they map. Join with *_OBJECTS_AE, *_TAB_COL, or both, for additional information.
*_EDITIONING_VIEW_COLS_AE	Describes the columns of every editioning view in the database, in every edition.

Each row of *_EDITIONING_VIEWS matches exactly one row of *_VIEWS, and each row of *_VIEWS that has EDITIONING_VIEW = 'Y' matches exactly one row of *_EDITIONING_VIEWS. Therefore, in this example, the WHERE clause is redundant:

```
SELECT ...
  FROM DBA_EDITIONING_VIEWS INNER JOIN DBA_VIEWS
  USING (OWNER, VIEW_NAME)
  WHERE EDITIONING_VIEW = 'Y'
  AND ...
```

The row of *_VIEWS that matches a row of *_EDITIONING_VIEWS has EDITIONING_VIEW = 'Y' by definition. Conversely, no row of *_VIEWS that has EDITIONING_VIEW = 'N' has a counterpart in *_ EDITIONING_VIEWS.

Using Edition-Based Redefinition to Upgrade an Application

To use edition-based redefinition to upgrade your application online, you must first ready your application:

1. Editions-enable the appropriate users.

   For instructions, see "Enabling Editions for a User" on page 19-4.

   The reason for this step is that only editions-enabled users can own editioning views, which you create in the next step.

2. Prepare your application to use editioning views.

   For instructions, see "Preparing Your Application to Use Editioning Views" on page 19-29.

With the editioning views in place, you can use edition-based redefinition to upgrade your application online as often as necessary. For each upgrade:

- If the type of every object that you will redefine is editionable (tables are not editionable), use the "Procedure for Edition-Based Redefinition Using Only Editions" on page 19-30.

- If you will change the structure of one or more tables, and while you are doing so, other users *do not* need to be able to change data in those tables, use the "Procedure for Edition-Based Redefinition Using Editioning Views" on page 19-32.
- If you will change the structure of one or more tables, and while you are doing so, other users must be able to change data in those tables, use the "Procedure for Edition-Based Redefinition Using Crossedition Triggers" on page 19-33.

Topics:

- Preparing Your Application to Use Editioning Views
- Procedure for Edition-Based Redefinition Using Only Editions
- Procedure for Edition-Based Redefinition Using Editioning Views
- Procedure for Edition-Based Redefinition Using Crossedition Triggers
- Rolling Back the Application Upgrade
- Reclaiming Space Occupied by Unused Table Columns
- Example: Using Edition-Based Redefinition to Upgrade an Application

Preparing Your Application to Use Editioning Views

An application that uses one or more tables must cover each table with an editioning view. An editioning view **covers** a table when all of these statements are true:

- Every ordinary object in the application references the table only through the editioning view. (An **ordinary object** is any object except an editioning view or crossedition trigger. Editioning views and crossedition triggers must reference tables.)
- Application users are granted object privileges only on the editioning view, not on the table.
- Oracle Virtual Private Database (VPD) policies are attached only to the editioning view, not to the table. (Regular auditing and fine-grained auditing (FGA) policies are attached only to the table.)

When the editioning view is actualized, a copy of the VPD policy is attached to the actualized editioning view. (A policy is uniquely identified by its name and the object to which is it attached.) If the policy function is also actualized, the copy of the policy uses the actualized policy function; otherwise, it uses the original policy function.

The static data dictionary views `*_POLICIES`, which describe the VPD policies, can have different results in different editions.

See Also:

- *Oracle Database Security Guide* for information about VPD, including that static data dictionary views that show information about VPD policies
- *Oracle Database Reference* for information about `*_POLICIES`

If an existing application does not use editioning views, prepare it to use them by following this procedure for each table that it uses:

1. Give the table a new name (so that you can give its current name to its editioning view).

Oracle recommends choosing a new name that is related to the original name and reflects the change history. For example, if the original table name is `Data`, the new table name might be `Data_1`.

2. (Optional) Give each column of the table a new name.

   Again, Oracle recommends choosing new names that are related to the original names and reflect the change history. For example, `Name` and `Number` might be changed to `Name_1` and `Number_1`.

   Any triggers that depend on renamed columns are now invalid. For details, see the entry for `ALTER TABLE` *`table`* `RENAME` *`column`* in Table 18–2.

3. Create the editioning view, giving it the original name of the table.

   For instructions, see "Creating an Editioning View" on page 19-16.

   Because the editioning view has the name that the table had, objects that reference that name now reference the editioning view.

4. If triggers are defined on the table, drop them, and rerun the code that created them.

   Now the triggers that were defined on the table are defined on the editioning view.

5. If VPD policies are attached to the table, drop the policies and policy functions and rerun the code that created them.

   Now the VPD policies that were attached to the table are attached to the editioning view.

6. Revoke all object privileges on the table from all application users.

   To see which application users have which object privileges on the table, use this query:

   ```
   SELECT GRANTEE, PRIVILEGE
   FROM DBA_TAB_PRIVS
   WHERE TABLE_NAME='table_name';
   ```

7. For every privilege revoked in step 6, grant the same privilege on the editioning view.

8. Enable editions for users who own private synonyms that refer to the table (for instructions, see "Enabling Editions for a User" on page 19-4) and notify those users that they must re-create those synonyms.

   Note: Public synonyms that refer to the table now fail with ORA-00980, and you cannot re-create them on the editioning view (for the reason, see "Actualizing Referenced Objects" on page 19-9).

Procedure for Edition-Based Redefinition Using Only Editions

Use this procedure only if the type of every object that you will redefine is editionable (as defined in "Editionable and Noneditionable Schema Object Types" on page 19-3). Table is not an editionable type.

1. Create a new edition.

   For instructions, see "Creating an Edition" on page 19-5.

2. Make the new edition your session edition.

   For instructions, see "Changing Your Session Edition" on page 19-11.

3. Make the necessary changes to the editioned objects of the application.
4. Ensure that all objects are valid.

   Query the static data dictionary `*_OBJECTS_AE`, which describes every actual object in the database, in every edition. If invalid objects remain, recompile them, using any `UTL_RECOMP` subprogram (described in *Oracle Database PL/SQL Packages and Types Reference*).
5. Check that the changes work as intended.

   If so, go to step 6.

   If not, either make further changes (return to step 3) or roll back the application upgrade (for instructions, see "Rolling Back the Application Upgrade" on page 19-34).
6. Make the new edition (the upgraded application) available to all users.

   For instructions, see "Making an Edition Available to All Users" on page 19-10.
7. Retire the old edition (the original application), so that all users except `SYS` use only the upgraded application.

   For instructions, see "Retiring an Edition" on page 19-13.

Example 19–6 shows how to use the preceding procedure to change a very simple PL/SQL procedure.

Example 19–6 Edition-Based Redefinition of Very Simple Procedure

1. Create PL/SQL procedure for this example:

```
CREATE OR REPLACE PROCEDURE hello IS
BEGIN
  DBMS_OUTPUT.PUT_LINE('Hello, edition 1.');
END hello;
/
```

2. Invoke PL/SQL procedure:

```
BEGIN hello(); END;
/
```

   Result:

```
Hello, edition 1.

PL/SQL procedure successfully completed.
```

3. Do edition-based redefinition of procedure:

   a. Create new edition:

```
CREATE EDITION e2 AS CHILD OF ora$base;
```

      Result:

```
Edition created.
```

   b. Make new edition your session edition:

```
ALTER SESSION SET EDITION = e2;
```

      Result:

```
Session altered.
```

c. Change procedure:

```
CREATE OR REPLACE PROCEDURE hello IS
BEGIN
  DBMS_OUTPUT.PUT_LINE('Hello, edition 2.');
END hello;
/
```

Result:

```
Procedure created.
```

d. Check that change works as intended:

```
BEGIN hello(); END;
/
```

Result:

```
Hello, edition 2.
PL/SQL procedure successfully completed.
```

e. Make new edition available to all users (requires system privileges):

```
ALTER DATABASE DEFAULT EDITION = e2;
```

f. Retire old edition (requires system privileges):

List grantees:

```
SELECT GRANTEE, PRIVILEGE
FROM DBA_TAB_PRIVS
WHERE TABLE_NAME = UPPER('ora$base')
/
```

Result:

```
GRANTEE                        PRIVILEGE
------------------------------ ---------
PUBLIC                         USE

1 row selected.
```

Revoke use on old edition from all grantees:

```
REVOKE USE ON EDITION ora$base FROM PUBLIC;
```

Procedure for Edition-Based Redefinition Using Editioning Views

Use this procedure only if you will change the structure of one or more tables, and while you are doing so, other users *do not* need to be able to change data in those tables.

1. Create a new edition.

   For instructions, see "Creating an Edition" on page 19-5.

2. Make the new edition your session edition.

   For instructions, see "Changing Your Session Edition" on page 19-11.

3. In the new edition, if the editioning views are read-only, make them read-write.

For instructions, see "Changing the 'Write-ability' of an Editioning View" on page 19-17.

4. In every edition except the new edition, make the editioning views read-only.
5. Make the necessary changes to the objects of the application.
6. Ensure that all objects are valid.

   Query the static data dictionary `*_OBJECTS_AE`, which describes every actual object in the database, in every edition. If invalid objects remain, recompile them, using any `UTL_RECOMP` subprogram (described in *Oracle Database PL/SQL Packages and Types Reference*).
7. Check that the changes work as intended.

   If so, go to step 8.

   If not, either make further changes (return to step 5) or roll back the application upgrade (for instructions, see "Rolling Back the Application Upgrade" on page 19-34).
8. Make the upgraded application available to all users.

   For instructions, see "Making an Edition Available to All Users" on page 19-10.
9. Retire the old edition (the original application), so that all users except `SYS` use only the upgraded application.

   For instructions, see "Retiring an Edition" on page 19-13.

Procedure for Edition-Based Redefinition Using Crossedition Triggers

Use this procedure only if you will change the structure of one or more tables, and while you are doing so, other users must be able to change data in those tables.

1. Create a new edition.

   For instructions, see "Creating an Edition" on page 19-5.
2. Make the new edition your session edition.

   For instructions, see "Changing Your Session Edition" on page 19-11.
3. Make the permanent changes to the objects of the application.

   For example, add new columns to the tables and create any new permanent subprograms.

   Objects that depend on objects that you changed might now be invalid. For more information, see Table 18–2.
4. Ensure that all objects are valid.

   Query the static data dictionary `*_OBJECTS_AE`, which describes every actual object in the database, in every edition. If invalid objects remain, recompile them, using any `UTL_RECOMP` subprogram (described in *Oracle Database PL/SQL Packages and Types Reference*).
5. Create the temporary objects—the crossedition triggers (in the disabled state) and any subprograms that they need.

   For instructions, see "Creating a Crossedition Trigger" on page 19-22.

   You need reverse crossedition triggers only if you do step 10, which is optional.
6. When the crossedition triggers compile successfully, enable them.

Use the `ALTER TRIGGER` statement with the `ENABLE` option. For information about this statement, see *Oracle Database PL/SQL Language Reference*.

7. Wait until pending changes are either committed or rolled back.

   Use the procedure `DBMS_UTILITY.WAIT_ON_PENDING_DML`, described in *Oracle Database PL/SQL Packages and Types Reference*.

8. Apply the transforms.

   For instructions, see "Transforming Data from Pre- to Post-Upgrade Representation" on page 19-24.

   Note: It is impossible to predict whether this step visits an existing row before a user of an ancestor edition updates, inserts, or deletes data from that row.

9. Check that the changes work as intended.

   If so, go to step 10.

   If not, either make further changes (return to step 3) or roll back the application upgrade (for instructions, see "Rolling Back the Application Upgrade" on page 19-34).

10. (Optional) Grant the `USE` privilege on your session edition to the early users of the upgraded application.

    For instructions, see "Making an Edition Available to Some Users" on page 19-10.

11. Make the upgraded application available to all users.

    For instructions, see "Making an Edition Available to All Users" on page 19-10.

12. Disable or drop the constraints and then drop the crossedition triggers.

    For instructions, see "Dropping the Crossedition Triggers" on page 19-26.

13. Retire the old edition (the original application), so that all users except `SYS` use only the upgraded application.

    For instructions, see "Retiring an Edition" on page 19-13.

Rolling Back the Application Upgrade

To roll back the application upgrade:

1. Change your session edition to something other than the new edition that you created for the upgrade.

   For instructions, see "Changing Your Session Edition" on page 19-11.

2. Drop the new edition that you created for the upgrade.

   For instructions, see "Dropping an Edition" on page 19-14.

3. If you created new table columns during the upgrade, reclaim the space that they occupy (for instructions, see "Reclaiming Space Occupied by Unused Table Columns" on page 19-34).

Reclaiming Space Occupied by Unused Table Columns

If you roll back an upgrade for which you created new table columns,

To reclaim the space that unused columns occupy:

1. Set the values of the unused columns to NULL.

   To avoid locking out other users while doing this operation, use the DBMS_PARALLEL_EXECUTE procedure (described in *Oracle Database PL/SQL Packages and Types Reference*).

2. Set the unused columns to UNUSED.

   Use the ALTER TABLE statement (described in *Oracle Database SQL Language Reference*) with the SET UNUSED clause (described in *Oracle Database SQL Language Reference*).

3. Shrink the table.

   Use the ALTER TABLE statement (described in *Oracle Database SQL Language Reference*) with the SHRINK SPACE clause (described in *Oracle Database SQL Language Reference*).

Example: Using Edition-Based Redefinition to Upgrade an Application

This example uses an edition, an editioning view, a forward crossedition trigger, and a reverse crossedition trigger.

Topics:

- Existing Application
- Preparing the Application to Use Editioning Views
- Using Edition-Based Redefinition to Upgrade the Application

> **Note:** Before you can use edition-based redefinition to upgrade an application, you must enable editions for every schema that the application uses. For instructions, see "Enabling Editions for a User" on page 19-4.

Existing Application

The existing application—the application to be upgraded—consists of a single table on which a trigger is defined. The application was created as in Example 19–7.

Example 19–7 Creating the Existing Application

1. Create table:

```
CREATE TABLE Contacts(
  ID            NUMBER(6,0) CONSTRAINT Contacts_PK PRIMARY KEY,
  Name          VARCHAR2(47),
  Phone_Number  VARCHAR2(20)
);
```

2. Populate table (not shown).

3. Prepare to create trigger on table:

```
ALTER TABLE Contacts ENABLE VALIDATE CONSTRAINT Contacts_PK;

DECLARE Max_ID INTEGER;
BEGIN
  SELECT MAX(ID) INTO Max_ID FROM Contacts;
```

```
  EXECUTE IMMEDIATE '
    CREATE SEQUENCE Contacts_Seq
      START WITH '||To_Char(Max_ID + 1);
END;
/
```

4. Create trigger:

```
CREATE TRIGGER Contacts_BI
  BEFORE INSERT ON Contacts FOR EACH ROW
BEGIN
  :NEW.ID := Contacts_Seq.NEXTVAL;
END;
/
```

Example 19–8 shows how the table `Contacts` looks after being populated with data.

Example 19–8 Viewing Data in Existing Table

Query:

```
SELECT * FROM Contacts
ORDER BY Name;
```

Result:

```
        ID NAME                                           PHONE_NUMBER
---------- ------------------------------------------------ --------------------
       174 Abel, Ellen                                      011.44.1644.429267
       166 Ande, Sundar                                     011.44.1346.629268
       130 Atkinson, Mozhe                                  650.124.6234
       105 Austin, David                                    590.423.4569
       204 Baer, Hermann                                    515.123.8888
       116 Baida, Shelli                                    515.127.4563
       167 Banda, Amit                                      011.44.1346.729268
       172 Bates, Elizabeth                                 011.44.1343.529268
       192 Bell, Sarah                                      650.501.1876
       151 Bernstein, David                                 011.44.1344.345268
       129 Bissot, Laura                                    650.124.5234
       169 Bloom, Harrison                                  011.44.1343.829268
       185 Bull, Alexis                                     650.509.2876
       187 Cabrio, Anthony                                  650.509.4876
       148 Cambrault, Gerald                                011.44.1344.619268
       154 Cambrault, Nanette                               011.44.1344.987668
       110 Chen, John                                       515.124.4269
       ...
       120 Weiss, Matthew                                   650.123.1234
       200 Whalen, Jennifer                                 515.123.4444
       149 Zlotkey, Eleni                                   011.44.1344.429018

107 rows selected.
```

Suppose that you must redefine `Contacts`, replacing the `Name` column with the columns `First_Name` and `Last_Name`, and adding the column `Country_Code`. Also suppose that while you are making this structural change, other users must be able to change the data in `Contacts`.

You need all features of edition-based redefinition: the edition, which is always needed; the editioning view, because you are redefining a table; and crossedition triggers, because other users must be able to change data in the table while you are redefining it.

Preparing the Application to Use Editioning Views

Example 19–9 shows how to create the editioning view from which other users will access the table `Contacts` while you are redefining it in the new edition.

Example 19–9 Creating an Editioning View for the Existing Table

1. Give table a new name (so that you can give its current name to editioning view):

```
ALTER TABLE Contacts RENAME TO Contacts_Table;
```

2. (Optional) Give columns of table new names:

```
ALTER TABLE Contacts_Table
  RENAME COLUMN Name TO Name_1;

ALTER TABLE Contacts_Table
  RENAME COLUMN Phone_Number TO Phone_Number_1;
```

3. Create editioning view:

```
CREATE OR REPLACE EDITIONING VIEW Contacts AS
  SELECT
    ID                 ID,
    Name_1             Name,
    Phone_Number_1     Phone_Number
  FROM Contacts_Table;
```

4. Move trigger `Contacts_BI` from table to editioning view:

```
DROP TRIGGER Contacts_BI;

CREATE TRIGGER Contacts_BI
  BEFORE INSERT ON Contacts FOR EACH ROW
BEGIN
  :NEW.ID := Contacts_Seq.NEXTVAL;
END;
/
```

Using Edition-Based Redefinition to Upgrade the Application

Example 19–10 shows how to create an edition in which to upgrade the "Existing Application" on page 19-35, make the new edition the session edition, and check that the new edition really is the session edition.

Example 19–10 Creating Edition in Which to Upgrade the Application

1. Create new edition:

```
CREATE EDITION Post_Upgrade AS CHILD OF Ora$Base;
```

2. Make new edition your session edition:

```
ALTER SESSION SET EDITION = Post_Upgrade;
```

3. Check session edition:

```
SELECT
SYS_CONTEXT('Userenv', 'Current_Edition_Name') "Current_Edition"
FROM DUAL;
```

Result:

```
Current_Edition
```

```
--------------------------------------------------------------------------------
POST_UPGRADE

1 row selected.
```

In the `Post_Upgrade` edition, Example 19–11 shows how to add the new columns to the physical table and recompile the trigger that was invalidated by adding the columns. Then, it shows how to replace the editioning view `Contacts` so that it selects the columns of the table by their desired logical names.

Example 19–11 Changing the Table and Replacing the Editioning View

1. Add new columns to physical table:

```
ALTER TABLE Contacts_Table ADD (
  First_Name_2     varchar2(20),
  Last_Name_2      varchar2(25),
  Country_Code_2   varchar2(20),
  Phone_Number_2   varchar2(20)
);
```

   (This is nonblocking DDL.)

2. Recompile invalidated trigger:

```
ALTER TRIGGER Contacts_BI COMPILE REUSE SETTINGS;
```

3. Replace editioning view so that it selects replacement columns with their desired logical names:

```
CREATE OR REPLACE EDITIONING VIEW Contacts AS
  SELECT
    ID                 ID,
    First_Name_2       First_Name,
    Last_Name_2        Last_Name,
    Country_Code_2     Country_Code,
    Phone_Number_2     Phone_Number
  FROM Contacts_Table;
```

In the `Post_Upgrade` edition, Example 19–12 shows how to create two procedures for the forward crossedition trigger to use, create both the forward and reverse crossedition triggers in the disabled state, and enable them.

Example 19–12 Creating and Enabling the Crossedition Triggers

1. Create first procedure that forward crossedition trigger uses:

```
CREATE OR REPLACE PROCEDURE Set_First_And_Last_Name (
  Name        IN  VARCHAR2,
  First_Name  OUT VARCHAR2,
  Last_Name   OUT VARCHAR2)
IS
  Comma_Pos NUMBER := INSTR(Name, ',');
BEGIN
  IF Comma_Pos IS NULL OR Comma_Pos < 2 THEN
    RAISE Program_Error;
  END IF;

  Last_Name := SUBSTR(Name, 1, Comma_Pos-1);
  Last_Name := RTRIM(Ltrim(Last_Name));

  First_Name := SUBSTR(Name, Comma_Pos+1);
```

```
  First_Name := RTRIM(LTRIM(First_Name));
END Set_First_And_Last_Name;
/
```

2. Create second procedure that forward crossedition trigger uses:

```
CREATE OR REPLACE PROCEDURE Set_Country_Code_And_Phone_No (
  Phone_Number     IN  VARCHAR2,
  Country_Code     OUT VARCHAR2,
  Phone_Number_V2  OUT VARCHAR2)
IS
  Char_To_Number_Error EXCEPTION;
  PRAGMA EXCEPTION_INIT(Char_To_Number_Error, -06502);
  Bad_Phone_Number EXCEPTION;
  Nmbr VARCHAR2(30) := REPLACE(Phone_Number, '.', '-');

  FUNCTION Is_US_Number(Nmbr IN VARCHAR2)
    RETURN BOOLEAN
  IS
    Len NUMBER := LENGTH(Nmbr);
    Dash_Pos NUMBER := INSTR(Nmbr, '-');
    n PLS_INTEGER;
  BEGIN
    IF Len IS NULL OR Len <> 12 THEN
      RETURN FALSE;
    END IF;
    IF Dash_Pos IS NULL OR Dash_Pos <> 4 THEN
      RETURN FALSE;
    END IF;
    BEGIN
      n := TO_NUMBER(SUBSTR(Nmbr, 1, 3));
    EXCEPTION WHEN Char_To_Number_Error THEN
      RETURN FALSE;
    END;

    Dash_Pos := INSTR(Nmbr, '-', 5);

    IF Dash_Pos IS NULL OR Dash_Pos <> 8 THEN
      RETURN FALSE;
    END IF;

    BEGIN
      n := TO_NUMBER(SUBSTR(Nmbr, 5, 3));
    EXCEPTION WHEN Char_To_Number_Error THEN
      RETURN FALSE;
    END;

    BEGIN
      n := TO_NUMBER(SUBSTR(Nmbr, 9));
    EXCEPTION WHEN Char_To_Number_Error THEN
      RETURN FALSE;
    END;

    RETURN TRUE;
  END Is_US_Number;

BEGIN
  IF Nmbr LIKE '011-%' THEN
    DECLARE
      Dash_Pos NUMBER := INSTR(Nmbr, '-', 5);
    BEGIN
```

```
      Country_Code := '+'|| TO_NUMBER(SUBSTR(Nmbr, 5, Dash_Pos-5));
      Phone_Number_V2 := SUBSTR(Nmbr, Dash_Pos+1);
    EXCEPTION WHEN Char_To_Number_Error THEN
      raise Bad_Phone_Number;
    END;
  ELSIF Is_US_Number(Nmbr) THEN
    Country_Code := '+1';
    Phone_Number_V2 := Nmbr;
  ELSE
    RAISE Bad_Phone_Number;
  END IF;
EXCEPTION WHEN Bad_Phone_Number THEN
  Country_Code := '+0';
  Phone_Number_V2 := '000-000-0000';
END Set_Country_Code_And_Phone_No;
/
```

3. Create forward crossedition trigger in disabled state:

```
CREATE OR REPLACE TRIGGER Contacts_Fwd_Xed
  BEFORE INSERT OR UPDATE ON Contacts_Table
  FOR EACH ROW
  FORWARD CROSSEDITION
  DISABLE
BEGIN
  Set_First_And_Last_Name(
    :NEW.Name_1,
    :NEW.First_Name_2,
    :NEW.Last_Name_2
  );
  Set_Country_Code_And_Phone_No(
    :NEW.Phone_Number_1,
    :NEW.Country_Code_2,
    :NEW.Phone_Number_2
  );
END Contacts_Fwd_Xed;
/
```

4. Enable forward crossedition trigger:

```
ALTER TRIGGER Contacts_Fwd_Xed ENABLE;
```

5. Create reverse crossedition trigger in disabled state:

```
CREATE OR REPLACE TRIGGER Contacts_Rvrs_Xed
  BEFORE INSERT OR UPDATE ON Contacts_Table
  FOR EACH ROW
  REVERSE CROSSEDITION
  DISABLE
BEGIN
  :NEW.Name_1 := :NEW.Last_Name_2||', '||:NEW.First_Name_2;
  :NEW.Phone_Number_1 :=
  CASE :New.Country_Code_2
    WHEN '+1' THEN
      REPLACE(:NEW.Phone_Number_2, '-', '.')
    ELSE
      '011.'||LTRIM(:NEW.Country_Code_2, '+')||'.'||
      REPLACE(:NEW.Phone_Number_2, '-', '.')
  END;
END Contacts_Rvrs_Xed;
/
```

6. Enable reverse crossedition trigger:

```
ALTER TRIGGER Contacts_Rvrs_Xed ENABLE;
```

7. Wait until pending changes are either committed or rolled back:

```
DECLARE
  scn              NUMBER  := NULL;
  timeout CONSTANT INTEGER := NULL;
BEGIN
  IF NOT DBMS_UTILITY.WAIT_ON_PENDING_DML(Tables  => 'Contacts_Table',
                                          timeout => timeout,
                                          scn     => scn)
  THEN
    RAISE_APPLICATION_ERROR(-20000,
     'Wait_On_Pending_DML() timed out. CETs were enabled before SCN: '||SCN);
  END IF;
END;
/
```

For information about the `DBMS_UTILITY.WAIT_ON_PENDING_DML` procedure, see *Oracle Database PL/SQL Packages and Types Reference*.

In the `Post_Upgrade` edition, Example 19–13 shows how to apply the transforms.

Example 19–13 Applying the Transforms

```
DECLARE
  c NUMBER := DBMS_SQL.OPEN_CURSOR();
  x NUMBER;
BEGIN
  DBMS_SQL.PARSE(
    c                          => c,
    Language_Flag              => DBMS_SQL.NATIVE,
    Statement                  => 'UPDATE Contacts_Table SET ID = ID',
    Apply_Crossedition_Trigger => 'Contacts_Fwd_Xed'
  );
  x := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.CLOSE_CURSOR(c);
  COMMIT;
END;
/
```

In the `Post_Upgrade` edition, Example 19–14 shows how to check that the change worked as intended. Compare Example 19–14 to Example 19–8.

Example 19–14 Viewing Data in Changed Table

1. Format columns for readability:

```
COLUMN ID FORMAT 999
COLUMN Last_Name FORMAT A15
COLUMN First_Name FORMAT A15
COLUMN Country_Code FORMAT A12
COLUMN Phone_Number FORMAT A12
```

2. Query:

```
SELECT * FROM Contacts
ORDER BY Last_Name;
```

Result:

```
  ID FIRST_NAME      LAST_NAME       COUNTRY_CODE PHONE_NUMBER
---- --------------- --------------- ------------ ------------
 174 Ellen           Abel            +44          1644-429267
 166 Sundar          Ande            +44          1346-629268
 130 Mozhe           Atkinson        +1           650-124-6234
 105 David           Austin          +1           590-423-4569
 204 Hermann         Baer            +1           515-123-8888
 116 Shelli          Baida           +1           515-127-4563
 167 Amit            Banda           +44          1346-729268
 172 Elizabeth       Bates           +44          1343-529268
 192 Sarah           Bell            +1           650-501-1876
 151 David           Bernstein       +44          1344-345268
 129 Laura           Bissot          +1           650-124-5234
 169 Harrison        Bloom           +44          1343-829268
 185 Alexis          Bull            +1           650-509-2876
 187 Anthony         Cabrio          +1           650-509-4876
 154 Nanette         Cambrault       +44          1344-987668
 148 Gerald          Cambrault       +44          1344-619268
 110 John            Chen            +1           515-124-4269
       ...
 120 Matthew         Weiss           +1           650-123-1234
 200 Jennifer        Whalen          +1           515-123-4444
 149 Eleni           Zlotkey         +44          1344-429018

107 rows selected.
```

If the change worked as intended, you can now follow steps 10 through 13 of the "Procedure for Edition-Based Redefinition Using Crossedition Triggers" on page 19-33.

A

Multithreaded extproc Agent

This appendix explains what the multithreaded `extproc` agent is, how it contributes to the overall efficiency of a distributed database system, and how to administer it.

Topics:

- Why Use the Multithreaded extproc Agent?
- Multithreaded extproc Agent Architecture
- Administering the Multithreaded extproc Agent

Why Use the Multithreaded extproc Agent?

This section explains how the multithreaded `extproc` agent contributes to the efficiency of external procedures.

Topics:

- The Challenge of Dedicated Agent Architecture
- The Advantage of Multithreading

The Challenge of Dedicated Agent Architecture

By default, an `extproc` agent is started for each user session and the `extproc` agent process terminates only when the user session ends.

This architecture can consume an unnecessarily large amount of system resources. For example, suppose that several thousand user sessions simultaneously spawn `extproc` agent processes. Because an `extproc` agent process is started for each session, several thousand `extproc` agent processes run concurrently. The `extproc` agent processes operate regardless of whether each individual `extproc` agent process is active at the moment. Thus `extproc` agent processes and open connections can consume a disproportionate amount of system resources. When sessions connect to Oracle Database, this problem is addressed by starting the server in shared server mode. Shared server mode allows database connections to be shared by a small number of server processes.

The Advantage of Multithreading

The Oracle Database shared server architecture assumes that even when several thousand user sessions are open, only a small percentage of these connections are active at any given time. In shared server mode, there is a pool of shared server processes. User sessions connect to dispatcher processes that place the requested tasks

in a queue. The tasks are picked up by the first available shared server processes. The number of shared server processes is usually less that the number of user sessions.

The multithreaded `extproc` agent provides similar functionality for connections to external procedures. The multithreaded `extproc` agent architecture uses a pool of shared agent threads. The tasks requested by the user sessions are put in a queue and are picked up by the first available multithreaded `extproc` agent thread. Because only a small percentage of user connections are active at a given moment, using a multithreaded `extproc` architecture allows more efficient use of system resources.

Multithreaded extproc Agent Architecture

One multithreaded `extproc` agent must be started for each system identifier (SID) before attempting to connect to the external procedure. This is done using the agent control utility `agtctl`. This utility is also used to configure the agent and to shut down the agent.

Each Oracle Net listener that is running on a system listens for incoming connection requests for a set of SIDs. If the SID in an incoming Oracle Net connect string is an SID for which the listener is listening, then that listener processes the connection. Further, if a multithreaded `extproc` agent was started for the SID, then the listener passes the request to that `extproc` agent.

In the architecture for multithreaded `extproc` agents, each incoming connection request is processed by different kinds of threads:

- A single **monitor** thread. The monitor thread is responsible for:
  - Maintaining communication with the listener
  - Monitoring the load on the process
  - Starting and stopping threads when required
- Several **dispatcher** threads. The dispatcher threads are responsible for:
  - Handling communication with the Oracle Database
  - Passing task requests to the task threads
- Several **task** threads. The task threads handle requests from the Oracle Database processes.

Figure A–1 illustrates the architecture of the multithreaded `extproc` agent. User sessions 1 and 2 issue requests for callouts to functions in some DLLs. These requests get serviced through heterogeneous services to the multithreaded `extproc` agent. These requests get handled by the agent's dispatcher threads, which then pass them on to the task threads. The task thread that is actually handling a request is responsible for loading the respective DLL and calling the function therein.

- All requests from a user session get handled by the same dispatcher thread. For example, dispatcher 1 handles communication with user session 1, and dispatcher 2 handles communication with user session 2. This is the case for the lifetime of the session.
- The individual requests can be serviced by different task threads. For example, task thread 1 can handle the request from user session 1, and later handle the request from user session 2.

> **See Also:** *Oracle Database Administrator's Guide*. for details on managing processes for external procedures

Figure A–1 Multithreaded extproc Agent Architecture

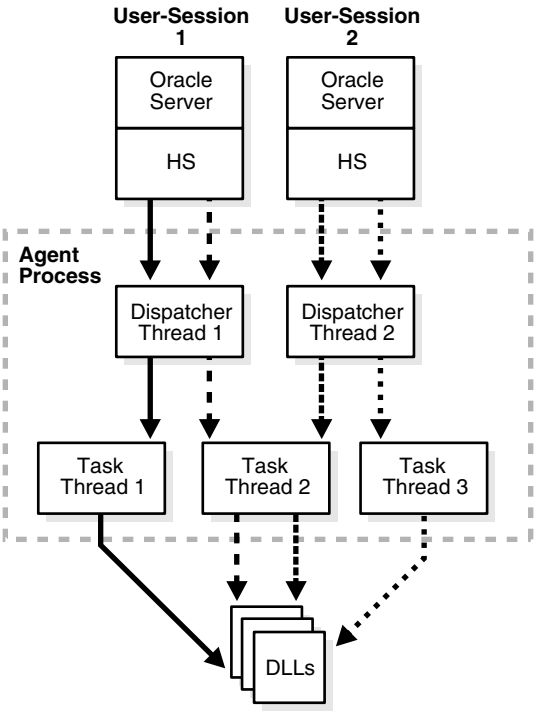

These three thread types roughly correspond to the Oracle Database multithreaded server PMON, dispatcher, and shared server processes, respectively.

Note: All requests from a user session go through the same dispatcher thread, but can be serviced by different task threads. Also, several task threads can use the same connection to the external procedure.

These topics explain each type of thread in more detail:

- Monitor Thread
- Dispatcher Threads
- Task Threads

See Also: "Administering the Multithreaded extproc Agent" on page A-4 for more information about starting and stopping the multithreaded `exproc` agent by using the agent control utility `agtctl`

Monitor Thread

When the agent control utility `agtctl` starts a multithreaded `extproc` agent for a SID, `agtctl` creates the monitor thread. The monitor thread performs these functions:

- Creates the dispatcher and task threads.
- Registers the dispatcher threads with all the listeners that are handling connections to this `extproc` agent. While the dispatcher for this SID is running, the listener does not start a process when it gets an incoming connection. Instead, the listener gives the connection to this same dispatcher.

- Monitors the other threads and sends load information about the dispatcher threads to all the listener processes handling connections to this `extproc` agent, enabling listeners to give incoming connections to the least loaded dispatcher.
- Continues to monitor each of the threads it has created.

Dispatcher Threads

Dispatcher threads perform these functions:

- Accept incoming connections and task requests from Oracle Database servers.
- Place incoming requests on a queue for a task thread to pick up.
- Send results of a request back to the server that issued the request.

> **Note:** After a user session establishes a connection with a dispatcher, all requests from that user session go to the same dispatcher until the end of the user session.

Task Threads

Task threads perform these functions:

- Pick up requests from a queue.
- Perform the necessary operations.
- Place the results on a queue for a dispatcher to pick up.

Administering the Multithreaded extproc Agent

One multithreaded `extproc` agent must be started for each system identifier (SID) before attempting to connect to the external procedure.

A multithreaded `extproc` agent is started, stopped, and configured by an agent control utility called `agtctl`, which works like `lsnrctl`. However, unlike `lsnrctl`, which reads a configuration file (`listener.ora`), `agtctl` takes configuration information from the command line and writes it to a control file.

Before starting `agtctl`, ensure that Oracle Listener is running. Then use the `agtctl` commands to set the `agtctl` configuration parameters (if you do not want their default values) and to start `agtctl`, as in Example A–1.

Example A–1 Setting Configuration Parameters and Starting agtctl

```
agtctl set max_dispatchers  2 ep_agt1
agtctl set tcp_dispatchers  1   ep_agt1
agtctl set max_task_threads 2  ep_agt1
agtctl set max_sessions     5 ep_agt1
agtctl unset listener_address ep_agt1
agtctl set listener_address "(address=(protocol=ipc)(key=extproc))" ep_agt1
agtctl startup extproc ep_agt1
```

You can use `agtctl` commands in either single-line command mode or shell mode.

Topics:

- Agent Control Utility (agtctl) Commands
- Using agtctl in Single-Line Command Mode

- Using Shell Mode Commands
- Configuration Parameters for Multithreaded extproc Agent Control

Agent Control Utility (agtctl) Commands

You can start and stop `agtctl` and create and maintain its control file by using the commands shown in Table A–1.

Table A–1 Agent Control Utility (agtctl) Commands

Command	Description
`startup`	Starts a multithreaded `extproc` agent
`shutdown`	Stops a multithreaded `extproc` agent
`set`	Sets a configuration parameter for a multithreaded `extproc` agent
`unset`	Causes a parameter to revert to its default value
`show`	Displays the value of a configuration parameter
`delete`	Deletes the entry for a particular SID from the control file
`exit`	Exits shell mode
`help`	Lists available commands

These commands can be issued in one of two ways:

- You can issue commands from the UNIX or DOS shell. This mode is called single-line command mode.
- You can enter `agtctl` and an `AGTCTL>` prompt appears. You then can enter commands from within the `agtctl` shell. This mode is called shell mode.

The syntax and parameters for `agtctl` commands depend on the mode in which they are issued.

Note:

- All commands are case-sensitive.
- The `agtctl` utility puts its control file in the directory specified by either one of two environment variables, `AGTCTL_ADMIN` or `TNS_ADMIN`. Ensure that at least one of these environment variables is set and that it specifies a directory to which the agent has access.
- If the multithreaded `extproc` agent requires that an environment variable be set, or if the `ENVS` parameter was used when configuring the `listener.ora` entry for the agent working in dedicated mode, then all required environment variables must be set in the UNIX or DOS shell that runs the `agtctl` utility.

Using agtctl in Single-Line Command Mode

This section describes the use of `agtctl` commands. They are presented in single-line command mode.

Setting Configuration Parameters for a Multithreaded extproc Agent

Set the configuration parameters for a multithreaded `extproc` agent before you start the agent. If a configuration parameter is not specifically set, a default value is used. Configuration parameters and their default values are shown in Table A–2.

Use the `set` command to set multithreaded `extproc` agent configuration parameters.

Syntax

```
agtctl set parameter parameter_value agent_sid
```

`parameter` is the parameter that you are setting.

`parameter_value` is the value being assigned to that parameter.

`agent_sid` is the SID that this agent services. This must be specified for single-line command mode.

Example

```
agtctl set max_dispatchers 5 salesDB
```

Starting a Multithreaded extproc Agent

Use the `startup` command to start a multithreaded `extproc` agent.

Syntax

```
agtctl startup extproc agent_sid
```

`agent_sid` is the SID that this multithreaded `extproc` agent services. This must be specified for single-line command mode.

Example

```
agtctl startup extproc salesDB
```

Shutting Down a Multithreaded extproc Agent

Use the `shutdown` command to stop a multithreaded `extproc` agent. There are three forms of shutdown:

- Normal (default)

  `agtctl` asks the multithreaded `extproc` agent to terminate itself gracefully. All sessions complete their current operations and then shut down.

- Immediate

  `agtctl` tells the multithreaded `extproc` agent to terminate immediately. The agent exits immediately regardless of the state of current sessions.

- Abort

  Without talking to the multithreaded `extproc` agent, `agtctl` issues a system call to stop it.

Syntax

```
agtctl shutdown [immediate|abort] agent_sid
```

`agent_sid` is the SID that the multithreaded `extproc` agent services. It must be specified for single-line command mode.

Example

```
agtctl shutdown immediate salesDB
```

Examining the Value of Configuration Parameters

To examine the value of a configuration parameter, use the `show` command.

Syntax

```
agtctl show parameter agent_sid
```

`parameter` is the parameter that you are examining.

`agent_sid` is the SID that this multithreaded `extproc` agent services. This must be specified for single-line command mode.

Example

```
agtctl show max_dispatchers salesDB
```

Resetting a Configuration Parameter to Its Default Value

You can reset a configuration parameter to its default value using the `unset` command.

Syntax

```
agtctl unset parameter agent_sid
```

`parameter` is the parameter that you are resetting (or changing).

`agent_sid` is the SID that this multithreaded `extproc` agent services. It must be specified for single-line command mode.

Example

```
agtctl unset max_dispatchers salesDB
```

Deleting an Entry for a Specific SID from the Control File

The `delete` command deletes the entry for the specified SID from the control file.

Syntax

```
agtctl delete agent_sid
```

`agent_sid` is the SID entry to delete.

Example

```
agtctl delete salesDB
```

Requesting Help

Use the `help` command to view a list of available commands for `agtctl` or to see the syntax for a particular command.

Syntax

```
agtctl help [command]
```

`command` is the name of the command whose syntax you want to view. The default is all `agtctl` commands.

Example

```
agtctl help set
```

Using Shell Mode Commands

In shell mode, start `agtctl` by entering:

```
agtctl
```

This results in the prompt `AGTCTL>`. Thereafter, because you are issuing commands from within the `agtctl` shell, you need not prefix the command string with `agtctl`.

Set the name of the agent SID by entering:

```
AGTCTL> set agent_sid agent_sid
```

All subsequent commands are assumed to be for the specified SID until the `agent_sid` value is changed. Unlike single-line command mode, you do not specify *`agent_sid`* in the command string.

You can set the language for error messages as follows:

```
AGTCTL> set language language
```

The commands themselves are the same as those for the single-line command mode. To exit shell mode, enter `exit`.

The following examples use shell mode commands.

Example: Setting a Configuration Parameter

This example sets a value for the `shutdown_address` configuration parameter.

```
AGTCTL> set shutdown_address (address=(protocol=ipc)(key=oraDBsalesDB))
```

Example: Starting a Multithreaded extproc Agent

This example starts a multithreaded `extproc` agent.

```
AGTCTL> startup extproc
```

Configuration Parameters for Multithreaded extproc Agent Control

Table A–2 describes and gives the defaults of the configuration parameters for the agent control utility.

Table A–2 Configuration Parameters for agtctl

Parameter	Description	Default Value
`max_dispatchers`	Maximum number of dispatchers	1
`tcp_dispatchers`	Number of dispatchers listening on TCP (the rest are using IPC)	0
`max_task_threads`	Maximum number of task threads	2

Table A–2 (Cont.) Configuration Parameters for agtctl

Parameter	Description	Default Value
`max_sessions`	Maximum number of sessions for each task thread	5

`listener_address`	Address on which the listener is listening (needed for registration)	`(ADDRESS_LIST=`
		`(ADDRESS=`
		`(PROTOCOL=IPC)`
		`(KEY=PNPKEY))`
		`(ADDRESS=`
		`(PROTOCOL=IPC)`
		`(KEY=`*`listener_sid`*`))`
		`(ADDRESS=`
		`(PROTOCOL=TCP)`
		`(HOST=127.0.0.1)`
		`(PORT=1521)))`
		Note: *`listener_sid`* is the IPC key of the address, on the Oracle Database, on which the listener is listening.

`shutdown_address`	Address the agent uses to communicate with the listener. This is the address on which the agent listens for all communication, including shutdown messages from `agtctl`.	`(ADDRESS=`
		`(PROTOCOL=IPC)`
		`(KEY=`*`listener_sid`* `||` *`agent_sid`*`))`
		`(ADDRESS=`
		`(PROTOCOL=TCP)`
		`(HOST=127.0.0.1)`
		`(PORT=1521))`
		Notes:
		▪ *`agent_sid`* is the SID of the multithreaded `extproc` agent.
		▪ `||` indicates that *`listener_sid`* and *`agent_sid`* are concatenated into one string.

max_dispatchers, tcp_dispatchers, max_task_threads, and max_sessions

To improve performance, you might need to change the values of some or all of the parameters `max_dispatchers`, `tcp_dispatchers`, `max_task_threads`, and `max_sessions`.

You can calculate the optimum values of `max_dispatchers`, `tcp_dispatchers`, `max_task_threads` with these formulas:

```
max_dispatchers = CEIL(x/y)
tcp_dispatchers = CEIL(x_tcpip/y)
max_task_threads = CEIL(x/max_sessions)
```

Where:

- `CEIL` is a SQL function that returns the smallest integer greater than or equal to its argument.
- *`x`* is the maximum number of sessions that can be connected to `extproc` concurrently.
- *`y`* is the maximum number of connections that the system can support for each dispatcher.
- *`x_tcpip`* is the maximum number of sessions that can be connected to `extproc` concurrently by TCP/IP.

  (*`x`* - *`x_tcpip`* is the maximum number of sessions that can be connected to `extproc` concurrently by IPC.)

There is no formula for computing the optimum value of `max_sessions`, which affects `max_task_threads`.

You must fine-tune these parameter settings, based on the capability of your hardware, and ensure that the concurrent threads do not exhaust your operating system.

The value of `max_dispatchers` must be at least 1 (which is the default).

Example

Suppose:

- The maximum number of sessions that can be connected to `extproc` concurrently (`x`) is 650.
- The maximum number of sessions that can be connected to `extproc` concurrently by TCP/IP (*`x_tcpip`*) is 400.

  (The maximum number of sessions that can be connected to `extproc` concurrently by IPC is 650-400=250.)
- The maximum number of connections that the system can support for each dispatcher (*`y`*) is 100.
- The maximum number of sessions for each task thread (`max_sessions`) is 20.

The optimum values for these parameters are:

```
max_dispatchers = CEIL(650/100) = CEIL(6.5)  = 7
tcp_dispatchers  = CEIL(400/100) = CEIL(4)    = 4
max_task_threads = CEIL(650/20)  = CEIL(32.5) = 33
```

That is, optimally:

- The maximum number of dispatchers is seven.
- Four of the seven dispatchers are listening on TCP/IP, and the remaining three are listening on IPC.
- The maximum number of task threads is 33.

listener_address and shutdown_address

The values of the configuration parameters `listener_address` and `shutdown_address` are specified with `ADDRESS`, as shown in both Table A–2 and Example A–1. Within `ADDRESS`, you can specify the parameter `HOST`, which can be either an IPv6 or IPv4 address or a host name. If `HOST` is a host name, then these values of the optional `ADDRESS` parameter `IP` are relevant:

IP Value	Meaning
`FIRST`	Listen on the first IP address returned by the DNS resolution of the host name.
`V4_ONLY`	Listen only on the IPv4 interfaces in the system.
`V6_ONLY`	Listen only on the IPv6 interfaces in the system.

For example, this value of `listener_address` or `shutdown_address` restricts it to IPv6 interfaces:

```
"(ADDRESS=(PROTOCOL=tcp)(HOST=sales)(PORT=1521)(IP=V6_ONLY))"
```

See Also: *Oracle Database Net Services Administrator's Guide* for detailed information about IPv6 support in Oracle Database

Index

Symbols

%ROWTYPE attribute, 6-6
%TYPE attribute, 6-6

Numerics

32-bit IEEE 754 format, 2-7
3GL (third-generation language), 6-33
64-bit IEEE 754 format, 2-7

A

Abstract Data Type (ADT)
  editions and, 19-3
  resetting evolved, 19-4
Active Server Pages (ASP), 10-3
actual object, 19-5
actualization
  in general, 19-5
  schema object dependency and, 19-9
address of row (rowid), 2-21
ADT
  editions and, 19-3
  resetting evolved, 19-4
AFTER SUSPEND trigger, 1-38
agent, 16-3
Agent Control Utility (agtctl)
  commands
    in shell mode, A-8
    in single-line mode, A-5
    list of, A-5
  extproc administration and, A-4
  extproc architecture and, A-2
aggregate function, 6-45
altering application online
  See edition-based redefinition
ancestor edition, 19-5
anonymous block, 6-2
ANSI/IDO data types, 2-19
AnyData data type, 2-17
AnyDataSet data type, 2-17
AnyType data type, 2-17
AP (application program), 15-4
application architecture, 13-2
application program (AP), 15-4
application SQL, 19-20
APPLYING_CROSSEDITION_TRIGGER
    function, 19-23
AQ (Oracle Advanced Queuing), 16-2
archive
  See Flashback Data Archive
ARG*n* data type, 2-27
arithmetic operation
  with date and time data type, 2-13
  with native floating-point data type, 2-8
ASP (Active Server Pages), 10-3
assignment
  data type conversion during, 2-24
  reported by PL/Scope, 7-9
attribute
  %ROWTYPE, 6-6
  %TYPE, 6-6
  Java STATIC class, 6-40
auditing policy, editioning view and, 19-29
Automatic Undo Management system, 12-1
autonomous transaction
  in general, 1-30
  trigger as, 1-38

B

backward compatibility
  LONG and LONG RAW data types for, 2-2
  RESTRICT_REFERENCES pragma for, 6-41
BATCH commit redo option, 1-6
BFILE data type, 2-15
binary format, 2-6
binary large object (BLOB) data type, 2-15
binary number, 2-5
binary ROWID, 2-24
BINARY_DOUBLE data type, 2-4
BINARY_FLOAT data type, 2-4
BLOB data type, 2-15
block, anonymous, 6-2
branch, 15-3
built-in function
  display type of, 2-27
  in regular expression, 3-2
  metadata for, 2-26
bulk binding
  in general, 6-16

when to use, 6-16

C

C external subprogram
callback with, 14-36
global variable in, 14-40
interface between PL/SQL and, 14-10
invoking, 14-29
loading, 14-4
passing parameter to, 14-15
publishing, 14-12
running, 14-26
service routine and, 14-30
See also external subprogram
C++ Class Library, 13-29
call specification
for external subprogram, 14-3
location of, 14-12
CALL statement, 14-26
calling subprogram
See invoking subprogram
cascading invalidation, 18-5
century, 2-12
CHANGE_DUPKEY_ERROR_INDEX hint, 19-23
CHAR data type
compared to VARCHAR2 data type, 2-4
specifying length of, 2-3
values stored, 2-2
character data type class, 18-16
character data types, 2-2
character large object (CLOB) data type, 2-15
CHECK constraint
compared to NOT NULL constraint, 5-14
designing, 5-14
multiple, 5-14
naming, 5-16
restrictions on, 5-13
when to use, 5-13
client notification, 16-2
client/server architecture, 13-2
CLOB data type, 2-15
cluster, creating index for, 4-2
coarse-grained invalidation, 18-5
collection
edition and, 19-3
referenced by DML statement, 6-16
referenced by FOR loop, 6-18
referenced by SELECT statement, 6-17
column
default value for
setting, 5-4
when to use, 5-4
multiple foreign key constraints on, 5-10
specifying length of, 2-3
commit redo management, 1-6
COMMIT statement, 1-5
committing transaction, 1-5
comparison operators, 2-8
compile-time error, handling
in general, 6-22
for multilanguage program, 14-30
composite FOREIGN KEY constraint, 5-7
composite index, 4-4
concurrency
serializable transaction for, 1-24
under explicit locking, 1-16
conditional expression represented as data, 2-20
connection pool, 13-15
constraint
in general, 5-1
altering, 5-20
CHECK
See CHECK constraint
crossedition trigger and
collisions, 19-23
dropping, 19-26
deferring checking of, 5-10
disabling
effect of, 5-16
existing, 5-19
new, 5-18
reasons for, 5-17
dropping, 5-22
editioning view and, 19-17
enabling
effect of, 5-16
existing, 5-18
new, 5-17
enforcing business rule with, 5-2
exception to, 5-19
FOREIGN KEY
See FOREIGN KEY constraint
minimizing overhead of, 5-12
naming, 5-16
NOT NULL
See NOT NULL constraint
on view, 5-6
PRIMARY KEY
See PRIMARY KEY constraint
privileges needed for defining, 5-16
renaming, 5-21
UNIQUE
See UNIQUE constraint
viewing definition of, 5-24
violating, 5-19
Continuous Query Notification (CQN), 11-1
converting data types
See data type conversion
copying on change, 19-5
coupling, 15-4
CQ_NOTIFICATION$_DESCRIPTOR object, 11-24
CQ_NOTIFICATION$_QUERY object, 11-26
CQ_NOTIFICATION$_REG_INFO object, 11-19
CQ_NOTIFICATION$_ROW object, 11-26
CQ_NOTIFICATION$_TABLE object, 11-25
CQN (Continuous Query Notification), 11-1
CREATE OR REPLACE optimization
in general, 18-5
actualization and, 19-5

crossedition trigger
in general, 19-18
creating, 19-22
displaying information about, 19-27
dropping, 19-26
execution of, 19-21
forward, 19-18
interaction with editions, 19-19
read-only editioning view and, 19-16
read-write editioning view and, 19-16
reverse, 19-19
sharing child cursor and, 19-27
crossedition trigger SQL
forward, 19-19
reverse, 19-20
cross-session PL/SQL function result cache, 6-10
current date and time, 2-11
current edition, 19-10
cursor
in general, 1-9
canceling, 1-11
closing, 1-10
crossedition trigger and, 19-27
number in session, 1-9
Oracle XA application and, 15-11
rerunning statement with, 1-10
schema object dependency and, 18-19
scrollable, 1-10
See also cursor variable
cursor variable
in general, 6-19
declaring, 6-19
examples of, 6-20
opening, 6-19

D

data definition language statement
See DDL statement
data integrity
See constraint
data type
in general, 2-2
ANSI/ISO, 2-19
DB2, 2-19
dynamic, 2-17
external, 2-2
family of, 2-26
for character data, 2-2
for date and time data, 2-11
for numeric data, 2-4
for specialized data, 2-14
object, 13-28
of formal subprogram parameter, 6-6
SQL/DS, 2-19
data type class, 18-16
data type conversion
in general, 2-24
of ANSI data types, 2-19
of ANSI/ISO data types, 2-19
of date and time data types, 2-13
of DB2 data types, 2-19
of native floating-point data types, 2-9
of SQL/DS data types, 2-19
date and time data types, 2-10
arithmetic operations with, 2-13
converting, 2-13
exporting, 2-14
importing, 2-14
DATE data type, 2-10
date data type class, 18-16
date format
changing default, 2-12
default, 2-12
DB2 data types, 2-19
DBMS_DEBUG package, 6-29
DBMS_DEBUG_JDWP package, 6-29
DBMS_FLASHBACK package, 12-14
DBMS_FLASHBACK.TRANSACTION_BACKOUT procedure, 12-15
DBMS_HPROF package, 8-2
DBMS_LOCK package, 1-22
DBMS_OUTPUT package, 6-27
DBMS_PARALLEL_EXECUTE package, 19-25
DBMS_STATS package, 12-26
DBMS_TYPES package, 2-17
DBMS_XA package, 15-16
DDL statement
Flashback Data Archive and, 12-22
ineffective, 18-5
Oracle XA and, 15-26
processing, 1-2
that generates notification, 11-5
deadlock, undetected, 1-22
debugging
PL/SQL Server Pages, 10-21
subprograms, 6-26
decimal number, 2-5
default column value
setting, 5-4
when to use, 5-4
default subprogram parameter value, 6-8
deferring constraint checks, 5-10
definer's-rights subprogram, 6-30
denormal floating-point number, 2-6
dependency mode, 18-12
dependent object
See schema object dependency
dependent transaction, 12-16
DESC keyword, 4-8
descendent edition, 19-5
DETERMINISTIC function
in general, 6-39
index and, 4-11
RPC signature and, 18-14
dirty read, 1-25
disabling constraint
effect of, 5-16
existing, 5-19
new, 5-18

reasons for, 5-17
dispatcher thread, A-2
distributed database
FOREIGN KEY constraint and, 5-12
remote dependency management and, 18-11
distributed query, run-time error in, 6-25
distributed transaction
in general, 15-3
how it works, 6-35
DLL (dynamic link library), 14-2
DML statement
bulk binding for, 6-16
parallel, 6-40
that references collection, 6-16
DML_LOCKS initialization parameter, 1-11
domain index, 4-7
double-precision IEEE 754 format, 2-7
drivers, Oracle JDBC, 13-6
DROP INDEX statement, 4-6
DTP (X/Open Distributed Transaction architecture), 15-1
dynamic link library (DLL), 14-2
dynamic registration, 15-4
dynamic SQL statement, 6-44
dynamically typed data, 2-17

E

edition
in general, 19-2
ancestor, 19-5
creating, 19-5
crossedition triggers and, 19-19
current, 19-10
descendent, 19-5
displaying information about, 19-27
enabling for a user, 19-4
leaf, 19-5
making available
to all users, 19-10
to some users, 19-10
ora$base, 19-2, 19-5
retiring, 19-13
root, 19-5
session, 19-10
visibility of trigger in, 19-19
editionable schema object type, 19-3
edition-based redefinition, 19-1
editioned object, 19-2
editioning view
in general, 19-15
auditing policy and, 19-29
changing base table of, 19-17
changing write-ability of, 19-17
covering table with, 19-29
creating, 19-16
displaying information about, 19-27
partition-extended name for, 19-16
preparing application for, 19-29
read-only, 19-16
read-write, 19-16
replacing, 19-17
SQL optimizer hint and, 19-18
editions-enabled user, 19-4
Electronic Product Code (EPC), 17-22
embedded PL/SQL gateway
in general, 9-3
how to use, 9-4
enabling constraint
effect of, 5-16
existing, 5-18
new, 5-17
enabling editions for a user, 19-4
encoding scheme, adding, 17-13
enforcing business rules
with application logic, 5-2
with constraints, 5-2
environment, programming, 13-1
EPC (Electronic Product Code), 17-22
error handling
compile-time
in general, 6-22
for multilanguage program, 14-30
run-time
See run-time error handling
exception
IEEE 754 standard
not raised, 2-8
raised during conversion, 2-9
in multilanguage program, 14-30
to constraint, 5-19
unhandled, 6-25
user-defined, 6-24
See also run-time error handling
EXCLUSIVE MODE option of LOCK TABLE statement, 1-14
EXPR data type, 2-27
expression
conditional, represented as data, 2-20
evaluation of, during data type conversion, 2-25
index built on
See function-based index
regular
See regular expression
expression directive in PSP script, 10-11
extended ROWID, 2-23
external binary ROWID, 2-24
external data type, 2-2
external large object (BFILE) data type, 2-15
external subprogram
in general, 14-2
call specification for, 14-3
debugging, 14-39
loading, 14-4
publishing, 14-9
external transaction manager, 15-3
extproc agent, A-1

F

families of data types, 2-26
fine-grained auditing (FGA) policy, editioning view and, 19-29
fine-grained invalidation, 18-5
firing order of triggers, 19-20
FIXED_DATE initialization parameter, 2-11
Flashback Data Archive, 12-18
Flashback Transaction, 12-15
FLASHBACK_TRANSACTION_QUERY view, 12-9
floating-point data type
  in general, 2-4
  range and precision of, 2-6
  See also native floating-point data type
floating-point number
  components of, 2-5
  denormal, 2-6
  format of, 2-5
  rounding, 2-5
  subnormal, 2-6
FOR loop
  bulk binding for, 6-18
  that references collection, 6-18
FORCE option of ALTER USER statement, 19-4
FOREIGN KEY constraint
  composite, 5-7
  distributed databases and, 5-12
  dropping, 5-22
  editioned view and, 19-3
  enabling, 5-22
  Flashback Transaction and, 12-16
  indexing, 5-12
  multiple, 5-10
  naming, 5-16
  NOT NULL constraint on, 5-9
  NULL value and, 5-8
  privileges needed to create, 5-23
  referential integrity enforced by, 5-23
  UNIQUE constraint on, 5-9
  without other constraints, 5-9
foreign key dependency, 12-16
foreign rowid, 2-24
formal subprogram parameter, data type of, 6-6
forward crossedition trigger, 19-18
forward crossedition trigger SQL, 19-19
function
  aggregate, 6-45
  built-in
    See built-in function
  controlling side effects of, 6-38
  DETERMINISTIC
    in general, 6-39
    index and, 4-11
    RPC signature and, 18-14
  invoking from SQL statement, 6-36
  MGD_ID ADT, 17-10
  OCI or OCCI, 13-19
  overloaded, 6-44
  PARALLEL_ENABLE, 6-39
    in general, 6-39
    RPC signature and, 18-14
  purity of
    in general, 6-38
    RPC signature and, 18-14
  result-cached, 6-10
  returning large amount of data from, 6-45
  SQL
    See built-in function
  See also subprogram
function result cache, 6-10
function-based index
  in general, 4-7
  editioned function and, 19-3
  examples of, 4-12
  optimizer and, 4-8
  sorting with, 4-8

G

Geographic Information System (GIS) data, 2-14
global transaction, 15-3
global variable, in C external subprogram, 14-40
group commit, 1-6

H

hierarchical profiler, 8-1
host language, 13-14
host program, 13-14
hot rollover, 19-1
HTML syntax error in PSP script, 10-7

I

IA-32 and IA-64 instruction set architecture, 2-9
IBM CICS, 15-3
IBM Transarc Encina, 15-3
Identity Code Package, 17-1
IEEE 754 format
  in general, 2-7
IEEE 754 standard
  data type
    See native floating-point data type
  exception
    not raised, 2-8
    raised during conversion, 2-9
  OCI support for datatypes of, 2-10
  special values supported by, 2-7
IGNORE_ROW_ON_DUPKEY_INDEX hint, 19-23
IMMEDIATE commit redo option, 1-6
IN OUT subprogram parameter mode, 6-6
IN subprogram parameter mode, 6-6
independent transaction
  See autonomous transaction
index
  application-specific, 4-1
  composite, 4-4
  creating
    examples of, 4-6
    for cluster, 4-2
    for table, 4-2

privileges needed for, 4-1
temporary table space for, 4-2
with SQL*Loader, 4-3
domain, 4-7
dropping, 4-5
edition-based redefinition and, 19-17
function-based
in general, 4-7
examples of, 4-12
on MGD_ID column, 17-10
overhead for, 4-4
statistics for, 4-5
where to put, 4-3
infinity, 2-7
inherited object
in general, 19-5
dropping, 19-7
initialization parameter
DML_LOCKS, 1-11
FIXED_DATE, 2-11
OPEN_CURSORS, 1-10
integer data type class, 18-16
integrity of data
See constraint
interface
between PL/SQL and C, 14-10
between PL/SQL and Java, 14-10
OraDatabase, 13-25
program, 13-3
TX, 15-4
user
in general, 13-3
stateful or stateless, 13-3
See also Oracle C++ Call Interface
See also Oracle Call Interface
INTERVAL DAY TO SECOND data type, 2-11
INTERVAL YEAR TO MONTH data type, 2-11
invalidation
cascading, 18-5
coarse-grained, 18-5
fine-grained, 18-5
of dependent object, 18-5
of package, 6-14
invoker's-rights subprogram, 6-30
invoking subprogram
in general, 6-29
from 3GL application, 6-33
from subprogram, 6-32
from trigger, 6-32
interactively from Oracle Database tools, 6-30
through embedded PL/SQL gateway, 9-17
isolation level
See transaction isolation level

J

Java class method
calling, 14-29
interface between PL/SQL and, 14-10
loading, 14-4
publishing, 14-11
See also external subprogram
Java Database Connectivity
See Oracle JDBC
Java language
compared to PL/SQL, 13-13
Oracle Database support for, 13-4
STATIC class attribute of, 6-40
Java Server Pages (JSP), 10-3
Java Virtual Machine
See Oracle JVM
JavaScript, 10-20
JDBC
See Oracle JDBC
JSP (Java Server Pages), 10-3
JVM
See Oracle JVM

K

key
foreign
See FOREIGN KEY constraint
primary
See PRIMARY KEY constraint
referential integrity
See FOREIGN KEY constraint
unique
See UNIQUE constraint

L

large object (LOB) data types
list of, 2-15
Oracle Objects for OLE support for, 13-28
leaf edition, 19-5
LGWR (log writer process), 1-6
libunit, 14-2
lightweight queue, 16-3
loadpsp utility, 10-13
LOB (large object) data types
list of, 2-15
Oracle Objects for OLE support for, 13-28
LOCK TABLE statement
in general, 1-12
SELECT FOR UPDATE statement with, 1-16
locking row explicitly, 1-15
locking table
explicitly, 1-11
implicitly, 1-14
log writer process (LGWR), 1-6
logical rowid, 2-24
LONG and LONG RAW data types, 2-15
loose coupling, 15-4
LOWER function, 4-8

M

main transaction, 1-30
maximum availability of table, 19-16
metacharacter in regular expression, 3-4

metadata for built-in function, 2-26
MGD_ID ADT, 17-1
MGD_ID database ADT function, 17-10
mod_plsql module, 9-2
mode
  agtctl command, A-4
  dependency, 18-12
  lock, 1-12
  serialized
    See serializable transaction
  subprogram parameter, 6-6
MODIFY CONSTRAINT clause of ALTER TABLE statement, 5-20
modifying
  See altering
monitor thread, A-2
multilanguage program
  in general, 14-1
  error or exception in, 14-30
multilingual extensions to POSIX standard, 3-7
multimedia data, 2-14
multithreaded extproc agent, A-1

N

name resolution
  in general, 18-10
  editions and, 19-9
NaN (not a number), 2-7
national character large object (NCLOB) data type, 2-15
native execution, compiling subprogram for, 6-18
native floating-point data type
  in general, 2-4
  arithmetic operation with, 2-8
  binary format for, 2-6
  clients that support, 2-10
  comparison operator for, 2-8
  conversion functions for, 2-9
  rounding, 2-5
  special values for, 2-7
NCHAR data type
  specifying length of, 2-3
  values stored, 2-2
NCLOB data type, 2-15
negative infinity, 2-7
negative zero, 2-7
new features, xxxi
NLSSORT function, 4-8
noneditionable schema object type, 19-3
noneditioned object, 19-2
nonpersistent queue, 16-3
normalized significand, 2-6
NOT NULL constraint
  compared to CHECK constraint, 5-14
  naming, 5-16
  on FOREIGN KEY constraint, 5-9
  when to use, 5-3
NOWAIT commit redo option, 1-6
NOWAIT option of LOCK TABLE statement, 1-12
NULL value and FOREIGN KEY constraint, 5-8
number
  binary, 2-5
  decimal, 2-5
  rounding, 2-5
NUMBER data type, 2-4
number data type class, 18-16
numeric data types, 2-4
NVARCHAR2 data type
  specifying length of, 2-3
  values stored, 2-2

O

object
  actual, 19-5
  dependent
    See schema object dependency
  editioned, 19-2
  inherited
    in general, 19-5
    dropping, 19-7
  large, 2-15
  noneditioned, 19-2
  potentially editioned, 19-2
  referenced
    See schema object dependency
  size limit for PL/SQL stored, 6-13
object change notification, 11-2
object data type, 13-28
object type
  See schema object type
OCCI
  See Oracle C++ Call Interface
OCI
  See Oracle Call Interface
ODC (Oracle Data Control), 13-29
ODP.NET, 13-21
online application upgrade
  See edition-based redefinition
OO4O
  See Oracle Objects for OLE
OPEN_CURSORS initialization parameter, 1-10
optimizer
  editioning view and, 19-18
  function-based index and, 4-8
  RPC signature and, 18-14
ora$base edition, 19-2, 19-5
ORA_ROWSCN pseudocolumn, 12-12
Oracle Advanced Queuing (AQ), 16-2
Oracle C++ Call Interface
  in general, 13-18
  building application with, 13-20
  kinds of functions in, 13-19
  procedural and nonprocedural elements of, 13-19
Oracle Call Interface
  in general, 13-18
  building application with, 13-20
  commit redo action in, 1-7
  compared to precompiler, 13-21

kinds of functions in, 13-19
procedural and nonprocedural elements of, 13-19
with Oracle XA, 15-12
Oracle Data Control (ODC), 13-29
Oracle Data Provider for .NET, 13-21
Oracle Database package
in general, 6-15
for writing low-level debugging code, 6-28
run-time error raised by, 6-23
Oracle Expression Filter, 2-20
Oracle Flashback Query, 12-6
Oracle Flashback Technology
in general, 12-1
application development features, 12-2
configuring database for, 12-3
database administration features, 12-3
performance guidelines for, 12-26
Oracle Flashback Transaction Query, 12-9
Oracle Flashback Version Query, 12-8
Oracle JDBC
in general, 13-5
compared to Oracle SQLJ, 13-10
sample program
2.0, 13-7
pre-2.0, 13-8
Oracle JDeveloper
in general, 6-27
Oracle SQLJ and, 13-10
Oracle JPublisher, 13-11
Oracle JVM, 13-5
Oracle Lock Management services, 1-22
Oracle Multimedia, 2-14
Oracle Objects for OLE
in general, 13-22
Automation Server, 13-23
C++ Class Library, 13-29
object data type support, 13-28
object model, 13-24
Oracle RAC and Oracle XA, 15-22
Oracle SQLJ
in general, 13-8
compared to Oracle JDBC, 13-10
Oracle JDeveloper and, 13-10
Oracle Text, 2-16
Oracle Total Recall, 12-18
Oracle Tuxedo, 15-3
Oracle Virtual Private Database (VPD) policy,
editioning view and, 19-29
Oracle XA
Oracle RAC and, 15-22
subprograms, 15-5
when to use, 15-1
OUT subprogram parameter mode, 6-6
out-of-space error, 1-37
overloaded subprogram, 6-11

P

package
in general, 6-10
advantages of, 6-10
body of, 6-11
creating, 6-13
invalidation of, 6-14
naming, 6-14
Oracle Database
in general, 6-15
for writing low-level debugging code, 6-28
run-time error raised by, 6-23
privileges needed to create, 6-14
privileges needed to drop, 6-14
session state and, 18-8
size limit for, 6-13
specification of, 6-11
synonym for, 6-34
package invalidation and, 6-14
package subprogram, 6-4
parallel DML statement, 6-40
parallel query, 6-40
PARALLEL_ENABLE function
in general, 6-39
RPC signature and, 18-14
parameter
initialization
See initialization parameter
subprogram
See subprogram parameter
parameter mode, 6-6
partition-extended editioning view name, 19-16
PERL-influenced extensions to POSIX standard, 3-8
persistent LOB instance, 2-15
persistent queue, 16-3
phantom read, 1-25
PL/Scope tool, 7-1
plshprof utility, 8-13
PL/SQL function result cache, 6-10
PL/SQL gateway, 9-2
PL/SQL hierarchical profiler, 8-1
PL/SQL language
in general, 13-4
compared to Java, 13-13
PL/SQL object
CREATE OR REPLACE and, 18-5
PL/SQL Server Pages
characteristics of, 10-5
elements of, 10-4
loading, 10-13
script error in, 10-7
PL/SQL unit
in general, 6-1
stored, 6-4
PL/SQL Web Toolkit, 9-3
pool, connection, 13-15
positive infinity, 2-7
positive zero, 2-7
POSIX metacharacter in regular expression, 3-4
POSIX standard
extensions to
multilingual, 3-7
PERL-influenced, 3-8

in regular expression, 3-4
potentially editioned object, 19-2
precompiler
in general, 13-14
compared to Oracle Call Interface, 13-21
Oracle XA and, 15-11
PRIMARY KEY constraint
in general, 5-5
dropping, 5-22
Flashback Transaction and, 12-16
naming, 5-16
primary key dependency, 12-16
privileges
for creating indexes, 4-1
for creating package, 6-9
for creating subprogram, 6-9
for debugging subprogram, 6-27
for defining constraint, 5-16
for dropping package, 6-10
for dropping packages, 6-14
for dropping subprogram, 6-10
for Oracle Flashback Technology, 12-5
for running subprogram, 6-30
revoked, object dependency and, 18-8
Pro*C/C++ precompiler
in general, 13-14
native floating-point data type support in, 2-10
Pro*COBOL precompiler, 13-16
procedure
PL/SQL Server Pages and, 10-9
See also subprogram
product code, 17-22
profiler, 8-1
program interface, 13-3
programming environment, 13-1
PSP
See PL/SQL Server Pages
public information, required, 15-4
publish-subscribe model, 16-1
purity of function
in general, 6-38
RPC signature and, 18-14

Q

quality-of-service flag, 11-20
query
parallel, 6-40
registering for Continuous Query Notification, 11-10
run-time error in distributed, 6-25
query result change notification, 11-2
queue, 16-3

R

Radio Frequency Identification (RFID) technology, 17-21
RAISE statement, 6-24
RAW data type, 2-15
raw data type class, 18-16
READ COMMITTED transaction isolation level
compared to SERIALIZABLE, 1-29
in Oracle Database, 1-25
transaction interactions with, 1-25
read consistency, 1-11
statement-level, 1-8
transaction-level
in general, 1-8
locking tables explicitly for, 1-11
read-only transaction for, 1-8
read lock, 1-27
READ UNCOMMITTED transaction isolation level
in Oracle Database, 1-25
transaction interactions with, 1-25
read-only editioning view, 19-16
read-only transaction, 1-8
read-write editioning view, 19-16
redefinition, edition-based, 19-1
redo information for transaction, 1-6
redo management, 1-6
referenced object
See schema object dependency
referential integrity
serializable transactions and, 1-27
trigger for enforcing, 1-28
referential integrity key
See FOREIGN KEY constraint
REGEXP_COUNT function, 3-3
REGEXP_INSTR function, 3-3
REGEXP_LIKE condition, 3-3
REGEXP_REPLACE function, 3-3
REGEXP_SUBSTR function, 3-3
registration
dynamic, 15-4
for Continuous Query Notification, 11-10
in publish-subscribe model, 16-3
static, 15-4
regular expression
in general, 3-1
built-in function in, 3-2
condition in, 3-2
in SQL statement, 3-10
metacharacter in, 3-4
Oracle implementation of, 3-2
POSIX standard support in, 3-4
relational operators, 2-8
remote dependency management, 18-11
remote procedure call dependency management, 18-12
repeatable read
in general, 1-8
read-only transaction for, 1-8
REPEATABLE READ transaction isolation level
in Oracle Database, 1-25
transaction interactions with, 1-25
repeatable reads
locking tables explicitly for, 1-11
required public information, 15-4
rerunning SQL statement, 1-10

resource manager (RM), 15-3
RESTRICT_REFERENCES pragma
for backward compatibility, 6-41
overloaded functions and, 6-44
static and dynamic SQL statements and, 6-44
restricted ROWID, 2-23
result cache, 6-10
resumable storage allocation, 1-37
RETENTION GUARANTEE clause for undo tablespace, 12-4
RETENTION option of ALTER TABLE statement, 12-5
reverse crossedition trigger, 19-19
reverse crossedition trigger SQL, 19-20
RFID (Radio Frequency Identification) technology, 17-21
RM (resource manager), 15-3
ROLLBACK statement, 1-7
rolling back transaction, 1-7
root edition, 19-5
roundiing floating-point numbers, 2-5
routine
See subprogram
row
address of (rowid), 2-21
locking explicitly, 1-15
ROW EXCLUSIVE MODE option of LOCK TABLE statement, 1-13
ROW SHARE MODE option of LOCK TABLE statement, 1-13
rowid
in general, 2-21
foreign, 2-24
logical, 2-24
universal (urowid), 2-24
ROWID data type, 2-23
ROWID pseudocolumn
in general, 2-22
CQN and, 11-12
See also rowid
ROWTYPE_MISMATCH exception, 6-21
RPC dependency management, 18-12
RPC-signature dependency mode, 18-13
RR date format element, 2-12
rule on queue, 16-3
rules engine, 16-4
run-time error handling
in general, 6-23
for distributed query, 6-25
for PL/SQL Server Pages (PSP) script, 10-7
for remote subprogram, 6-26
for storage allocation error, 1-37
for user-defined exception, 6-24
See also exception

S

SAVEPOINT statement, 1-8
schema object dependency
in general, 18-1
actualization and, 19-9
in distributed database, 18-11
invalidation and, 18-5
on nonexistence of other objects, 18-11
revoked privileges and, 18-8
shared pool and, 18-19
schema object type
editionable, 19-3
noneditionable, 19-3
scrollable cursor, 1-10
searchable text, 2-16
SELECT FOR UPDATE statement
in general, 1-15
LOCK TABLE statement with, 1-16
referential integrity and
inside trigger, 1-28
outside trigger, 1-28
SELECT statement
bulk binding for, 6-17
referencing collection with, 6-17
with AS OF clause, 12-6
with FOR UPDATE clause
See SELECT FOR UPDATE statement
with VERSIONS BETWEEN clause, 12-8
semi-available table, 19-16
serendipitous change, 19-24
data transformation collisions and, 19-23
identifying, 19-23
serializable transaction
for concurrency control, 1-24
interaction with, 1-25
referential integrity and, 1-27
SERIALIZABLE transaction isolation level
compared to READ COMMITTED, 1-30
in Oracle Database, 1-25
transaction interactions with, 1-25
See also serializable transaction
server-side programming, 13-2
service routine, C external subprogram and, 14-30
session edition, 19-10
session state, 18-8
session variable, 6-31
SET CONSTRAINTS statement, 5-10
SET TRANSACTION statement
with READ ONLY option, 1-8
SHARE MODE option of LOCK TABLE statement, 1-13
SHARE ROW EXCLUSIVE MODE option of LOCK TABLE statement, 1-14
shared SQL area, 1-3
side effects of function, controlling, 6-38
signature checking, 18-11
single-precision IEEE 754 format, 2-7
spatial data, 2-14
SQL area, shared, 1-3
SQL data type
See data type
SQL function
See built-in function
SQL optimizer hint and editioning view, 19-18

SQL statement
application, 19-20
crossedition trigger
forward, 19-19
reverse, 19-20
dynamic, 6-44
invoking PL/SQL function from, 6-36
processing
DDL statement, 1-2
stages of, 1-1
system management statement, 1-2, 1-3
rerunning, 1-10
static, 6-44
SQL*Loader, creating index with, 4-3
SQL/DS data types, 2-19
SQLJ
See Oracle SQLJ
SQLT_BDOUBLE data type, 2-10
SQLT_BFLOAT data type, 2-10
standalone subprogram, 6-4
state
session, 18-8
user interface and, 13-3
web application and, 9-25
statement
See SQL statement
statement-level read consistency, 1-8, 1-11
static registration, 15-4
static SQL statement, 6-44
static variable, in C external subprogram, 14-40
statistics
for application, 8-1
for identifier, 7-1
for index, 4-5
storage allocation error, 1-37
stored PL/SQL unit, 6-4
subnormal floating-point number, 2-6
subprogram
compiling for native execution, 6-18
creating, 6-8
definer's-rights, 6-30
editioned, 19-3
exception-handling, 6-24
external
See external subprogram
invoker's-rights, 6-30
invoking
See invoking subprogram
naming, 6-5
Oracle XA, 15-5
overloaded, 6-11
package, 6-4
parameter of
See subprogram parameter
privileges needed to debug, 6-27
privileges needed to run, 6-30
remote, 6-26
size limit for, 6-13
standalone, 6-4
synonym for, 6-34
See also function *and* procedure
subprogram parameter
in general, 6-5
composite variable as, 6-8
data type of formal, 6-6
default value of, 6-8
mode of, 6-6
subscriber, 16-3
subscription services, 16-4
synonym
CREATE OR REPLACE and, 18-5
for package, 6-34
for subprogram, 6-34
public, for editioned object, 19-3
SYSDATE function, 2-11
system management statement, 1-2, 1-3

T

table
locking
choosing strategy for, 1-12
explicitly, 1-11
implicitly, 1-14
with maximum availability, 19-16
with semi-availability, 19-16
Tag Data Translation Markup Language Schema, 17-4
task thread, A-2
temporary LOB instance, 2-15
temporary table space, 4-2
thin client configuration, 13-2
third-generation language (3GL), 6-33
thread
dispatcher, A-2
monitor, A-2
Oracle XA library, 15-15
task, A-2
three-tier architecture, 13-2
tight coupling, 15-4
time and date date types, 2-10
time data types
See date and time data types
time stamp checking, 18-11
time, changing default format of, 2-12
TIMESTAMP data type, 2-11
time-stamp dependency mode, 18-12
TIMESTAMP WITH LOCAL TIME ZONE data type, 2-11
TIMESTAMP WITH TIME ZONE data type, 2-11
TM (transaction manager), 15-3
TPM (transaction processing monitor), 15-3
transaction
autonomous
in general, 1-30
trigger as, 1-38
choosing isolation level for, 1-29
committing, 1-5
dependent, 12-16
distributed

in general, 15-3
how it works, 6-35
global, 15-3
grouping operations into, 1-4
improving performance of, 1-4
main, 1-30
read-only, 1-8
redo entry for, 1-6
rolling back, 1-7
savepoints for, 1-8
serializable
See serializable transaction
statements in, 1-8
transaction interaction
kinds of, 1-25
serializable transaction and, 1-25
transaction isolation level and, 1-25
transaction isolation level
choosing, 1-29
setting, 1-26
transaction interaction and, 1-25
transaction manager (TM), 15-3
transaction processing monitor (TPM), 15-3
transaction set consistency, 1-29
transaction-level read consistency
in general, 1-8
locking tables explicitly for, 1-11
read-only transaction for, 1-8
transform
in general, 19-18
applying, 19-25
trigger
in general, 6-18
AFTER SUSPEND, 1-38
as autonomous transaction, 1-38
crossedition
See crossedition trigger
enforcing referential integrity with, 1-28
in edition
firing order of, 19-20
visibility of, 19-19
what kind can fire, 19-19
invoking subprogram from, 6-32
size limit for, 6-13
TRUST keyword in RESTRICT_REFERENCES pragma, 6-43
two-phase commit protocol, 15-3
two-tier architecture, 13-2
TX interface, 15-4
type attribute, 6-6

U

undetected deadlock, 1-22
undo data, 12-1
UNDO_RETENTION parameter, 1-9
undoing transaction, 1-7
unhandled exception, 6-25
Unicode character literal, 2-3
UNIQUE constraint
crossedition trigger and, 19-23
dropping, 5-22
naming, 5-16
on FOREIGN KEY constraint, 5-9
when to use, 5-5
universal rowid (urowid), 2-24
unrepeatable read, 1-25
upgrading applications online
See edition-based redefinition
UPPER function, 4-8
UROWID data type, 2-24
user interface
in general, 13-3
stateful and stateless, 13-3
user lock, 1-22
user-defined exception, 6-24
user-defined type, as subprogram parameter, 6-8
UTLLOCKT.SQL script, 1-24

V

VARCHAR data type class, 18-16
VARCHAR2 data type
compared to CHAR data type, 2-4
specifying length of, 2-3
values stored, 2-2
variable
cursor
See cursor variable
in C external subprogram
global, 14-40
static, 14-40
VERSIONS_ENDSCN pseudocolumn, 12-8
VERSIONS_ENDTIME pseudocolumn, 12-8
VERSIONS_OPERATION pseudocolumn, 12-9
VERSIONS_STARTSCN pseudocolumn, 12-8
VERSIONS_STARTTIME pseudocolumn, 12-8
VERSIONS_XID pseudocolumn, 12-8
view
constraint on, 5-6
editioned
FOREIGN KEY constraint and, 19-3
materialized view and, 19-3
editioning
See editioning view
VPD policy, editioning view and, 19-29

W

WAIT commit redo option, 1-6
WAIT option of LOCK TABLE statement, 1-12
web application
in general, 9-1
implementing, 9-2
state and, 9-25
web page
See also PL/SQL Server Pages
web services, 13-12
web toolkit
See PL/SQL Web Toolkit

WORK option of ROLLBACK statement, 1-7
wrap utility, debugging and, 6-28
write-ability of editioning view, 19-17
write-after-write dependency, 12-16

X

xa_open string, 15-8
XML data, 2-16
XMLType data type, 2-16
X/Open Distributed Transaction architecture, 15-1
X/Open Distributed Transaction Processing (DTP) architecture, 15-1

Y

YY date format element, 2-12