

Oracle® SQL Developer

User's Guide

Release 17.2

E88161-03

June 2017

Provides conceptual and usage information about Oracle SQL Developer, a graphical tool that enables you to browse, create, edit, and delete (drop) database objects; run SQL statements and scripts; edit and debug PL/SQL code; manipulate and export data; migrate third-party databases to Oracle; view metadata and data in MySQL and third-party databases; and view and create reports.

Oracle SQL Developer User's Guide, Release 17.2

E88161-03

Copyright © 2006, 2017, Oracle and/or its affiliates. All rights reserved.

Primary Author: Celin Cherian

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware 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 that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

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

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	xxvii
Audience	xxvii
Documentation Accessibility	xxvii
Product Accessibility	xxvii
Related Documents.....	xxvii
Conventions.....	xxviii
Third-Party License Information	xxviii
Antlr 3.1.3	xxviii
Batik Version 1.7.....	xxviii
Bouncy Castle Crypto API 1.45	xxxix
JGoodies Looks 1.3.1	xxxix
JXLayer 4.0.....	xxxix
Log4j 1.2.9	xxxix
POI 2.5.1.....	xxxix
Regular Expression Package 1.3.....	xxxix
1 SQL Developer Concepts and Usage	
1.1 About SQL Developer	1-2
1.2 Installing and Getting Started with SQL Developer.....	1-2
1.3 SQL Developer User Interface	1-3
1.3.1 Menus for SQL Developer.....	1-7
1.3.2 Restoring the Original "Look and Feel"	1-12
1.4 Database Objects	1-13
1.4.1 Applications (Application Express 3.0.1 and Later).....	1-14
1.4.2 Cache Groups (Oracle TimesTen In-Memory Database).....	1-14
1.4.3 Chains.....	1-14
1.4.4 Credentials.....	1-15
1.4.5 Database Destinations	1-15
1.4.6 Database Links (Public and Private).....	1-15
1.4.7 Destination Groups	1-15
1.4.8 Directories.....	1-16
1.4.9 Editions	1-16

1.4.10	File Watchers	1-16
1.4.11	Functions	1-16
1.4.12	Indexes	1-17
1.4.13	Java Sources	1-18
1.4.14	Jobs	1-18
1.4.15	Job Classes	1-18
1.4.16	Operators	1-18
1.4.17	Materialized Views	1-19
1.4.18	Materialized View Logs	1-19
1.4.19	Multitenant Container Database (CDB)	1-19
1.4.20	Packages	1-20
1.4.21	Procedures	1-20
1.4.22	Programs	1-21
1.4.23	Queues	1-21
1.4.24	Queue Tables	1-21
1.4.25	Recycle Bin	1-22
1.4.26	Replication Schemes (Oracle TimesTen In-Memory Database)	1-22
1.4.27	Schedules	1-22
1.4.28	Sequences	1-22
1.4.29	Synonyms (Public and Private)	1-23
1.4.30	Tables	1-23
1.4.31	Triggers	1-25
1.4.32	Types	1-26
1.4.33	Users (Other Users)	1-26
1.4.34	Views	1-26
1.4.35	Window Groups	1-27
1.4.36	Windows	1-27
1.4.37	XML DB Repository	1-27
1.4.38	XML Schemas	1-28
1.4.39	Captured and Converted Database Objects (for Migration)	1-28
1.5	Database Connections	1-28
1.5.1	Using Folders to Group Connections	1-31
1.5.2	Sharing of Connections	1-31
1.5.3	Advanced Security for JDBC Connection to the Database	1-31
1.5.4	Connections with Operating System (OS) Authentication	1-32
1.5.5	Connections with Proxy Authentication	1-32
1.5.6	Connections with SSH Authentication	1-33
1.6	Entering and Modifying Data	1-34
1.7	Running and Debugging Functions and Procedures	1-36
1.7.1	Using Bookmarks When Editing Functions and Procedures	1-39
1.7.2	Remote Debugging	1-39
1.7.3	Displaying SQL Trace (.trc) Files	1-40
1.7.4	Using the PL/SQL Hierarchical Profiler	1-40

1.7.5	Setting Expression Watches	1-40
1.8	Using the SQL Worksheet	1-41
1.8.1	SQL*Plus Statements Supported and Not Supported in SQL Worksheet	1-44
1.8.2	Script Runner	1-45
1.8.3	Execution Plan	1-46
1.8.4	Autotrace Pane.....	1-46
1.8.5	DBMS Output Pane.....	1-46
1.8.6	OWA Output Pane	1-47
1.8.7	SQL History.....	1-47
1.8.8	Query Builder.....	1-47
1.8.9	Command-Line Interface for SQL Formatting.....	1-48
1.8.10	SQL Worksheet "Hints" for Formatting Output	1-48
1.8.11	Gauges: In the SQL Worksheet and User-Defined Reports	1-49
1.8.12	Entering OLAP DML Statements in the SQL Worksheet.....	1-50
1.9	Using Snippets to Insert Code Fragments.....	1-50
1.9.1	User-Defined Snippets.....	1-51
1.10	Finding Database Objects	1-52
1.11	Using Recent Objects.....	1-53
1.12	Using Versioning	1-53
1.12.1	About Subversion and SQL Developer.....	1-53
1.12.2	Pending Changes.....	1-54
1.13	Using DBA Features in SQL Developer	1-54
1.13.1	Container Database (CDB).....	1-57
1.13.2	Database Configuration.....	1-57
1.13.3	Database Status.....	1-58
1.13.4	Data Guard	1-59
1.13.5	Data Pump.....	1-59
1.13.6	Oracle Sharding.....	1-60
1.13.7	Performance	1-60
1.13.8	RMAN Backup/Recovery.....	1-62
1.13.9	Resource Manager.....	1-64
1.13.10	SQL Translator FrameWork.....	1-64
1.13.11	Scheduler	1-64
1.13.12	Security	1-65
1.13.13	Storage.....	1-65
1.13.14	Tuning.....	1-66
1.14	Scheduling Jobs Using SQL Developer	1-67
1.14.1	Scheduler Design Editor.....	1-68
1.15	Using the Cart.....	1-68
1.15.1	Cart Versus Database Export Wizard.....	1-71
1.16	Spatial Support in SQL Developer	1-71
1.16.1	Context Menu Operations on Spatial Data and Metadata	1-72
1.16.2	Map Visualization of Spatial Data	1-72

1.17	Change Manager Support in SQL Developer	1-74
1.17.1	Change Plan Administrators and Developers	1-75
1.17.2	Developing Change Plans with SQL Developer	1-76
1.17.3	Change Manager Actions to Perform with Enterprise Manager.....	1-78
1.18	Apache Hadoop Connectors Support in SQL Developer	1-79
1.18.1	Copy to Hadoop: Overview and Requirements	1-80
1.18.2	Oracle Loader for Hadoop: Overview and Requirements	1-81
1.18.3	Oracle SQL Connector for HDFS: Overview and Requirements	1-82
1.19	SQL Developer Reports	1-82
1.19.1	Bind Variables for Reports	1-84
1.19.2	About Your Database reports.....	1-84
1.19.3	All Objects reports.....	1-84
1.19.4	Application Express reports	1-85
1.19.5	ASH and AWR reports	1-85
1.19.6	Charts reports	1-85
1.19.7	Database Administration reports.....	1-85
1.19.8	Data Dictionary reports	1-85
1.19.9	Scheduler reports.....	1-86
1.19.10	PL/SQL reports	1-86
1.19.11	Security reports.....	1-86
1.19.12	Streams reports	1-87
1.19.13	Table reports	1-87
1.19.14	XML reports	1-89
1.19.15	Data Modeler reports.....	1-89
1.19.16	User Defined reports.....	1-89
1.20	SQL Developer Preferences.....	1-92
1.20.1	Environment.....	1-92
1.20.2	Change Management Parameters	1-93
1.20.3	Code Editor	1-94
1.20.4	Compare and Merge	1-98
1.20.5	Database.....	1-99
1.20.6	Debugger	1-109
1.20.7	External Editor	1-110
1.20.8	File Types.....	1-111
1.20.9	Global Ignore List.....	1-111
1.20.10	Migration	1-111
1.20.11	Mouse Actions	1-115
1.20.12	Shortcut Keys (Accelerator Keys)	1-115
1.20.13	SSH (Secure Shell)	1-116
1.20.14	Unit Test Parameters.....	1-116
1.20.15	Usage Reporting	1-116
1.20.16	Versioning	1-116
1.20.17	Web Browser and Proxy.....	1-121

1.20.18	XML Schemas.....	1-122
1.21	Location of User-Related Information.....	1-122
1.22	Java Development Kit (JDK) Location.....	1-123
1.23	Command-Line Interface for SQL Developer.....	1-123
1.24	Data Modeler in SQL Developer.....	1-124
1.25	Oracle TimesTen In-Memory Database Support.....	1-125
1.26	Using the Help.....	1-125
1.27	SQL Developer Usage Suggestions.....	1-126
1.27.1	SQL History Shortcuts.....	1-126
1.27.2	Unshared Worksheets.....	1-126
1.27.3	SQL Worksheet Bookmarks.....	1-126
1.27.4	Oracle Data Miner.....	1-126
1.27.5	Formatted Display of SQL Trace (.trc) Files.....	1-126
1.27.6	Keyboard Navigation: Alt + Page Down/Up to Move Among Tabs.....	1-126
1.27.7	Folders for Organizing Connections.....	1-126
1.27.8	Third-Party Databases and SQL Developer.....	1-127
1.27.9	Debugger Ports and Firewalls.....	1-127
1.27.10	Viewing Multiple Tables.....	1-127
1.27.11	Customizing SQL Developer Appearance.....	1-127
1.27.12	Maximizing Tab Panes.....	1-127
1.27.13	Default Path for Running Scripts.....	1-127
1.27.14	Shutting Down and Restarting the Database.....	1-127
1.27.15	Feature Requests.....	1-127
1.27.16	Discussion Forums.....	1-127
1.27.17	Help Text Font Size.....	1-128
1.27.18	Procedure and Function Signatures.....	1-128
1.27.19	Type-Ahead in Navigators.....	1-128
1.27.20	Extended Paste.....	1-128
1.27.21	Closing Tabbed Windows Using the Mouse Wheel.....	1-128
1.27.22	Go to Last Edit Location.....	1-128
1.27.23	Closing Tabbed Windows Using the Context Menu.....	1-128
1.27.24	List of All Open Windows.....	1-128
1.27.25	Go to Subprogram Implementation from Package Window.....	1-129
1.27.26	Select Multiple Table or Column Names in Completion Insight.....	1-129
1.27.27	Startup Time and Automatic Check for Updates.....	1-129
1.28	For More Information About SQL Developer.....	1-129

2 SQL Developer: Migrating Third-Party Databases

2.1	Migration: Basic Options and Steps.....	2-1
2.1.1	Migrating Using the Migration Wizard.....	2-2
2.1.2	Copying Selected Tables to Oracle.....	2-5
2.2	Migration: Background Information and Guidelines.....	2-6
2.2.1	Overview of Migration.....	2-6

2.2.2	Preparing a Migration Plan.....	2-7
2.2.3	Before You Start Migrating: General Information.....	2-12
2.2.4	Before You Start Migrating: Source-Specific Information.....	2-14
2.2.5	Capturing the Source Database.....	2-17
2.2.6	Creating and Customizing the Converted Model.....	2-18
2.2.7	Generating the DDL for the Oracle Schema Objects.....	2-19
2.2.8	Migrating the Data.....	2-19
2.2.9	Making Queries Case Insensitive.....	2-24
2.2.10	Testing the Oracle Database.....	2-25
2.2.11	Deploying the Oracle Database.....	2-29
2.3	SQL Developer User Interface for Migration.....	2-32
2.3.1	Migration Submenu.....	2-33
2.3.2	Other Menus: Migration Items.....	2-33
2.3.3	Migration Preferences.....	2-33
2.3.4	Migration Log Panes.....	2-33
2.3.5	Using the Translation Scratch Editor.....	2-34
2.4	Command-Line Interface for Migration.....	2-34

3 SQL Developer: Oracle REST Data Services Support

3.1	Installing Oracle REST Data Services.....	3-1
3.2	Oracle REST Data Services Administration.....	3-1
3.2.1	About Oracle REST Data Services.....	3-2
3.2.2	Oracle REST Data Services Administration Toolbar and Context Menu.....	3-2
3.2.3	Connecting to Oracle REST Data Services.....	3-3
3.2.4	REST Data Services Navigator: Global Settings and Database Settings.....	3-4
3.2.5	REST Data Services Navigator: NoSQL Store Configuration.....	3-10
3.2.6	REST Data Services Navigator: Reports.....	3-11
3.3	Automatically Enabling REST Access to a Schema, Table, or View (AutoREST).....	3-11
3.4	REST Development.....	3-12
3.4.1	About RESTful Services.....	3-12
3.4.2	RESTful Services Terminology.....	3-12
3.4.3	RESTful Services Requirements and Setup.....	3-13
3.4.4	REST Data Services in the Connections Navigator.....	3-14
3.4.5	REST Development Pane.....	3-14

4 SQL Developer: Unit Testing

4.1	Overview of Unit Testing.....	4-1
4.2	SQL Developer User Interface for Unit Testing.....	4-3
4.2.1	Unit Test Submenu.....	4-3
4.2.2	Other Menus: Unit Test Items.....	4-4
4.2.3	Unit Test Preferences.....	4-4
4.3	Unit Test Repository.....	4-4
4.3.1	Managing Repository Users and Administrators.....	4-5

4.4	Editing and Running a Unit Test.....	4-5
4.5	Using a Dynamic Value Query to Create a Unit Test.....	4-6
4.6	Using Lookups to Simplify Unit Test Creation	4-8
4.6.1	Providing Values for Input Fields.....	4-9
4.6.2	Automatically Creating Implementations	4-9
4.7	Using Variable Substitution in Validation Actions.....	4-10
4.8	Unit Test Library	4-11
4.9	Unit Test Reports	4-12
4.10	Exporting and Importing Unit Test Objects.....	4-12
4.11	Command-Line Interface for Unit Testing.....	4-13
4.12	Unit Testing Best Practices	4-14
4.12.1	Strategy	4-14
4.12.2	Test Suites.....	4-14
4.12.3	Test Naming	4-14
4.12.4	Avoiding Test Naming Clashes	4-15
4.12.5	Test Implementations	4-15
4.12.6	Library.....	4-15
4.12.7	Lookups	4-16
4.12.8	Test and Suite Execution	4-16
4.12.9	Reports	4-16
4.13	Example of Unit Testing (Tutorial)	4-16
4.13.1	Create the EMPLOYEES Table	4-17
4.13.2	Create the AWARD_BONUS Procedure	4-18
4.13.3	Create the Unit Testing Repository	4-18
4.13.4	Create a Unit Test	4-19
4.13.5	Run the Unit Test.....	4-20
4.13.6	Create and Run an Exception Unit Test.....	4-21
4.13.7	Create a Unit Test Suite	4-22
4.13.8	Run the Unit Test Suite.....	4-23

5 SQL Developer Tutorial: Creating Objects for a Small Database

5.1	Create a Table (BOOKS).....	5-2
5.2	Create a Table (PATRONS)	5-4
5.3	Create a Table (TRANSACTIONS)	5-5
5.4	Create a Sequence	5-7
5.5	Insert Data into the Tables.....	5-7
5.6	Create a View	5-8
5.7	Create a PL/SQL Procedure.....	5-9
5.8	Debug a PL/SQL Procedure	5-10
5.9	Use the SQL Worksheet for Queries	5-12
5.10	Script for Creating and Using the Library Tutorial Objects	5-13

6 SQL Developer Dialog Boxes and Wizards

6.1 Add Extension	6-1
6.2 Add Schema Error	6-2
6.3 Advanced Properties (Connections)	6-2
6.3.1 Proxy tab	6-2
6.4 Application Migration.....	6-2
6.5 Associate Repository	6-3
6.6 Cart Error (Objects Not Available).....	6-3
6.7 Change Type.....	6-3
6.8 Check for Updates	6-3
6.9 Choose Directory	6-4
6.10 Clone PDB to Oracle Cloud	6-4
6.11 Clone Pluggable Database.....	6-5
6.12 Color Palette and Custom Colors.....	6-5
6.13 Configure Extension.....	6-5
6.14 Configure File Type Associations.....	6-5
6.15 Copy Columns.....	6-6
6.16 Copy Objects.....	6-6
6.17 Copy to Hadoop / Append to Hadoop Table.....	6-6
6.18 Copy to Oracle.....	6-8
6.19 Component Palette	6-8
6.20 Component Palette: Configure Component Palette	6-8
6.21 Component Palette: Create Palette Page	6-9
6.22 Component Palette: New/Edit Code Snippet	6-9
6.23 Component Palette: New Section/Rename Section	6-9
6.24 Confirm Drop Application	6-9
6.25 Confirm Running SQL	6-9
6.26 Connection Has Uncommitted Changes	6-9
6.27 Create New Object.....	6-10
6.28 Create/Edit Chain	6-10
6.29 Create/Edit Credential	6-10
6.30 Create/Edit/Select Database Connection.....	6-11
6.30.1 Oracle tab.....	6-12
6.30.2 TimesTen tab	6-14
6.30.3 Amazon Redshift tab	6-14
6.30.4 DB2 tab.....	6-15
6.30.5 Hive tab.....	6-15
6.30.6 JDBC tab.....	6-16
6.30.7 MySQL tab.....	6-16
6.30.8 PostgreSQL tab	6-16
6.30.9 SQL Server and Sybase tab	6-17
6.30.10 Teradata tab.....	6-17

6.30.11	Creating and Editing Connections.....	6-17
6.31	New/Edit Cloud Connection.....	6-18
6.31.1	Advanced Connection Information.....	6-18
6.32	New/Edit NoSQL Connection.....	6-19
6.33	Rename Model (Migration).....	6-19
6.34	Delete Confirmation.....	6-19
6.35	Delete Confirmation (Migration).....	6-19
6.36	Rename Database Item (Migration).....	6-20
6.37	Select Connection.....	6-20
6.38	Connection Information.....	6-20
6.39	No Connection Found.....	6-20
6.40	Connection Rename Error.....	6-20
6.41	New Folder (Connections).....	6-21
6.42	Continue After Pause.....	6-21
6.43	Select Library.....	6-21
6.44	Create Library.....	6-21
6.45	CVS: Check Out from CVS.....	6-21
6.46	CVS: Create/Edit CVS Connection.....	6-22
6.47	CVS: Import to CVS.....	6-23
6.48	CVS: Log In to CVS.....	6-24
6.49	Data Import Wizard (Load Data).....	6-24
6.49.1	Data Preview.....	6-24
6.49.2	Import Method.....	6-24
6.49.3	Choose Columns.....	6-25
6.49.4	Column Definition.....	6-25
6.49.5	Finish.....	6-25
6.50	Export/Import Connections.....	6-26
6.50.1	Export Connections.....	6-26
6.50.2	Import Connections.....	6-26
6.51	Create/Edit Database Destination.....	6-27
6.52	Create/Edit Destination Group (Database or External).....	6-27
6.53	Create/Edit Database Link.....	6-28
6.54	Create/Edit Index.....	6-28
6.54.1	Definition.....	6-29
6.54.2	Properties.....	6-29
6.54.3	Storage.....	6-30
6.54.4	Partitions.....	6-30
6.54.5	DDL.....	6-30
6.55	Create File Watcher.....	6-30
6.56	Create Filter.....	6-31
6.57	Create Function or Procedure.....	6-31
6.58	Create/Edit Job.....	6-32
6.58.1	Job Details.....	6-32

6.58.2	Destination	6-33
6.58.3	Job Arguments	6-33
6.58.4	Notification.....	6-33
6.58.5	Properties.....	6-33
6.58.6	Summary/SQL	6-34
6.59	Create/Edit Job Class.....	6-34
6.60	Create/Edit Materialized View Log	6-35
6.61	Create PL/SQL Package	6-36
6.62	Create Pluggable Database.....	6-36
6.63	Create/Edit Program	6-37
6.64	Create/Edit Role	6-38
6.65	Create/Edit Schedule	6-39
6.66	Create/Edit Sequence	6-40
6.67	Create SQL File.....	6-40
6.68	Create/Edit Synonym	6-41
6.69	Create Table (quick creation)	6-41
6.70	Create/Edit Table (with advanced options).....	6-42
6.70.1	Columns pane	6-43
6.70.2	Constraints	6-44
6.70.3	Primary Key Constraint.....	6-44
6.70.4	Unique Constraints	6-45
6.70.5	Foreign Keys Constraints	6-46
6.70.6	Check Constraints	6-46
6.70.7	Indexes pane.....	6-47
6.70.8	In-Memory pane	6-47
6.70.9	Storage pane	6-48
6.70.10	Column Sequences pane	6-48
6.70.11	Table Properties pane	6-49
6.70.12	LOB Parameters pane	6-49
6.70.13	Partitioning pane	6-50
6.70.14	Subpartition Templates pane.....	6-51
6.70.15	Partition Definitions pane	6-51
6.70.16	Index Organized Properties pane	6-52
6.70.17	External Table Properties pane.....	6-52
6.70.18	Comment pane.....	6-53
6.70.19	DDL pane.....	6-53
6.70.20	Hive Tables.....	6-53
6.71	Storage Options.....	6-54
6.72	Create/Edit Tablespace	6-55
6.72.1	File Specifications tab.....	6-56
6.72.2	Properties tab	6-56
6.72.3	Default Parameters tab	6-57
6.72.4	DDL tab.....	6-57

6.73	Create Trigger.....	6-57
6.74	Create Type (User-Defined)	6-59
6.75	Create/Edit User.....	6-59
6.75.1	User tab	6-59
6.75.2	Granted Roles tab	6-59
6.75.3	System Privileges tab	6-60
6.75.4	Quotas tab.....	6-60
6.75.5	SQL tab.....	6-60
6.76	Create/Edit User Defined Report	6-60
6.77	Create/Edit User Defined Report Folder	6-62
6.78	Create/Edit View.....	6-62
6.78.1	SQL Query tab or pane	6-62
6.78.2	View Properties or Materialized View Properties pane	6-63
6.78.3	DDL tab or pane	6-65
6.78.4	Quick-Pick Objects pane.....	6-65
6.78.5	FROM Clause pane	6-66
6.78.6	SELECT Clause pane.....	6-66
6.78.7	WHERE Clause pane	6-67
6.78.8	GROUP BY Clause pane.....	6-67
6.78.9	HAVING Clause pane	6-67
6.78.10	ORDER BY Clause pane	6-67
6.79	Create Window	6-68
6.80	Create Window Group.....	6-69
6.81	Create XML Document from XML Schema	6-69
6.82	Create XML Schema	6-70
6.83	Custom Filters	6-70
6.84	Customize Toolbars.....	6-70
6.85	Data Pump Export	6-70
6.85.1	Source page	6-71
6.85.2	Tablespaces page	6-71
6.85.3	Schemas page	6-71
6.85.4	Tables page.....	6-71
6.85.5	Filter page	6-71
6.85.6	Table Data page	6-71
6.85.7	Options page	6-71
6.85.8	Output Files page	6-71
6.85.9	Job Schedule page.....	6-72
6.85.10	Summary page	6-72
6.86	Data Pump Import.....	6-72
6.86.1	Type page	6-72
6.86.2	Tablespaces page	6-72
6.86.3	Schemas page	6-72
6.86.4	Tables page.....	6-72

6.86.5	Filter page	6-73
6.86.6	Table Data page	6-73
6.86.7	Remapping page.....	6-73
6.86.8	Options page	6-73
6.86.9	Job Schedule page.....	6-73
6.86.10	Summary page	6-73
6.87	Database Copy	6-73
6.87.1	Source/Destination page	6-73
6.87.2	Object Types page	6-74
6.87.3	Select Schemas page.....	6-74
6.87.4	Select Tablespace page.....	6-74
6.87.5	Specify Objects page	6-74
6.87.6	Specify Data page	6-74
6.87.7	Summary page	6-74
6.88	Database Differences	6-75
6.88.1	Source/Destination page	6-75
6.88.2	Types to Diff page	6-76
6.88.3	Specify Objects page	6-76
6.88.4	Summary page.....	6-76
6.89	Database Export (Unload Database Objects and Data).....	6-77
6.89.1	Source/Destination page	6-77
6.89.2	Types to Export page	6-78
6.89.3	Specify Objects page	6-79
6.89.4	Specify Data page	6-79
6.89.5	Export Summary page	6-79
6.90	DDL Panel for Creating or Editing an Object.....	6-79
6.91	Debugger - Attach to JPDA	6-79
6.92	Deploy or Import Application	6-79
6.93	Deploy Objects to Cloud.....	6-80
6.93.1	Cloud.....	6-80
6.93.2	Authorization.....	6-81
6.93.3	Transfer	6-81
6.93.4	Deploy DDL	6-81
6.93.5	Deploy Data.....	6-81
6.93.6	Unsupported Cart Object for Cloud Deployment.....	6-81
6.93.7	Other Errors (Invalid or Missing Entries).....	6-82
6.94	Describe Object Window	6-82
6.95	Diff Objects	6-82
6.96	Drop Pluggable Database.....	6-84
6.97	Drop REST Services for Schema	6-84
6.98	Edit/View Value (Data or Table Column).....	6-84
6.99	Enter Bind Values	6-85
6.100	Erase from Disk.....	6-85

6.101	Error Writing to Export File	6-85
6.102	Export Error	6-85
6.103	Export Data	6-85
6.104	Export Objects	6-85
6.105	Export Resource Module, Privilege, and Role.....	6-86
6.106	Export: Unique Names Warning	6-87
6.107	External Locator Configuration	6-87
6.108	External Tools.....	6-87
6.109	Create/Edit External Tool	6-88
6.109.1	Program Options	6-88
6.109.2	Display	6-88
6.109.3	Integration	6-88
6.109.4	Availability	6-89
6.110	Choose Offline Options.....	6-89
6.111	Edit Font Location.....	6-89
6.112	Edit Header or Footer (PDF Formatting)	6-90
6.113	Edit Join	6-90
6.114	Feature Missing	6-90
6.115	Feature Required.....	6-90
6.116	File Processing	6-91
6.117	Filter	6-91
6.118	Filter Object Types	6-92
6.119	Filter Schemas.....	6-92
6.120	Filter Error.....	6-92
6.121	Find/Highlight	6-92
6.122	Git: Add.....	6-93
6.123	Git: Add All	6-93
6.124	Git: Add to .gitignore File.....	6-93
6.125	Git: Checkout Revision	6-93
6.126	Git: Clone from Git	6-94
6.127	Git: Commit	6-95
6.128	Git: Commit All.....	6-95
6.129	Git: Create Branch.....	6-96
6.130	Git: Create Tag.....	6-96
6.131	Git: Export Committed Changes	6-96
6.132	Git: Export Uncommitted Changes	6-97
6.133	Git: Fetch from Git	6-97
6.134	Git: Initialize Repository	6-98
6.135	Git: Merge	6-98
6.136	Git: Pull from Git	6-98
6.137	Git: Push to Git	6-99
6.138	Git: Revert	6-100
6.139	Insert Macro.....	6-100

6.140	Externally Modified Files.....	6-100
6.141	Find/Replace Text	6-101
6.142	Find Result	6-101
6.143	Format Properties (Advanced Formatting).....	6-101
6.144	Generate Patch	6-102
6.145	Go to Bookmark	6-102
6.146	Go to Line Number.....	6-102
6.147	Go to Line Number: Error	6-102
6.148	History.....	6-102
6.149	Import Using Oracle Loader for Hadoop.....	6-103
6.150	Import Using Oracle SQL Connector for HDFS.....	6-104
6.151	Install/Run Oracle REST Data Services	6-106
6.151.1	ORDS File Locations	6-106
6.151.2	Database Connection	6-106
6.151.3	Database Options	6-107
6.151.4	Select Tablespace	6-107
6.151.5	PL/SQL Gateway	6-107
6.151.6	APEX RESTful Services	6-107
6.151.7	Run Standalone Mode	6-107
6.151.8	ORDS Users.....	6-108
6.151.9	Summary.....	6-108
6.152	Install SQL Translator	6-108
6.152.1	Installing a Translator and Creating a Profile: Usage Notes.....	6-109
6.153	Invalid Database Version.....	6-110
6.154	Load Extension.....	6-110
6.155	Load Keyboard Scheme	6-110
6.156	LOB Descriptors.....	6-110
6.157	Maintain Spatial Metadata	6-111
6.158	Manage Columns.....	6-111
6.159	Manage Connections (REST).....	6-112
6.160	Manage Features and Updates	6-112
6.161	Map Connection.....	6-113
6.162	Merge Migration Projects	6-113
6.163	Modify Pluggable State.....	6-114
6.164	Modify Value.....	6-114
6.165	Data Move Details	6-115
6.166	New/Edit Change Set.....	6-115
6.167	New Procedure (Refactoring)	6-116
6.168	New/Edit Local or Remote Port Forward	6-116
6.169	Edit SSH Connection	6-116
6.170	No Lock	6-116
6.171	No Object Found.....	6-117
6.172	No Object Selected	6-117

6.173	Object Preferences.....	6-117
6.174	Offline Generation Error - Destination Directory	6-117
6.175	Open File	6-117
6.176	Oracle-Only Report	6-117
6.177	Paste	6-117
6.178	Perforce: Add Files to Perforce	6-117
6.179	Perforce: Connection	6-118
6.180	Perforce: Create Changelist.....	6-118
6.181	Perforce: Delete Files	6-118
6.182	Perforce: Login	6-119
6.183	Perforce: Open Files for Edit	6-119
6.184	Perforce: Revert Files.....	6-119
6.185	Perforce: Sync Files	6-119
6.186	Perforce: Submit Changelist.....	6-120
6.187	Plug In Pluggable Database	6-120
6.188	Print Preview	6-121
6.189	Privilege Warning for Migration.....	6-122
6.190	Publish to REST.....	6-122
6.191	Recent Files	6-122
6.192	Delete or Truncate Repository	6-122
6.192.1	Error Truncating Repository.....	6-123
6.193	Rename Local Variable	6-123
6.194	Rename Tab	6-123
6.195	Rename Procedure.....	6-123
6.196	Repair (Validate) Oracle REST Data Services	6-123
6.197	REST Data Services Connection	6-124
6.198	REST Data Services Connection: Authentication Failed	6-124
6.199	REST Data Services Connection: Remove Connection.....	6-126
6.200	RESTful Services Wizard (Auto-Enable REST Access)	6-126
6.200.1	Specify Details.....	6-126
6.200.2	RESTful Summary	6-127
6.201	RESTful Services Wizard and Object Dialog Boxes	6-127
6.201.1	Resource Module.....	6-128
6.201.2	Resource Template	6-128
6.201.3	Resource Handler	6-129
6.201.4	RESTful Summary	6-130
6.202	RESTful Services Connection	6-130
6.203	RESTful Services Privilege.....	6-131
6.204	RESTful Services Roles.....	6-131
6.205	Select Current Repository	6-132
6.206	Cannot Capture Table	6-132
6.207	Report Panel	6-132
6.208	Reset Expired Password (Enter New Password)	6-132

6.209	Revision Lister	6-132
6.210	Run/Debug/Profile PL/SQL	6-132
6.211	Create/Edit Breakpoint	6-133
6.211.1	Definition tab	6-133
6.211.2	Conditions tab.....	6-133
6.211.3	Actions tab.....	6-134
6.212	Save/Save As, or Select File	6-134
6.213	Save or Open Cart Tool Configuration.....	6-134
6.214	Save Files	6-134
6.215	Unable to Save Files.....	6-135
6.216	Save Style Settings	6-135
6.217	Schema Differences Source or Destination Error	6-135
6.218	Script Execution Failed	6-135
6.219	Script Generation Complete	6-135
6.220	Set Data Mapping	6-135
6.221	Add/Edit Rule	6-136
6.222	Set Encoding	6-136
6.223	Set Pause Continue	6-136
6.224	Shared Wizard Pages	6-137
6.224.1	Types	6-137
6.224.2	Specify Objects.....	6-137
6.224.3	Data.....	6-137
6.224.4	Summary.....	6-138
6.225	Sign In (checking for updates)	6-138
6.226	Single Record View.....	6-138
6.227	Save Snippet (User-Defined).....	6-138
6.228	Edit Snippets (User-Defined)	6-139
6.229	Show SQL.....	6-139
6.230	Start Date, End Date, Repeat Interval, Comments (Scheduler)	6-139
6.231	Subversion: Add Property	6-140
6.232	Subversion: Add to Source Control	6-140
6.233	Subversion: Apply Patch	6-140
6.234	Subversion: Branch/Tag.....	6-140
6.235	Subversion: Check Out from Subversion	6-141
6.236	Subversion: Commit Resources	6-141
6.237	Subversion: Commit Working Copy.....	6-142
6.238	Subversion: Confirm Checkout.....	6-142
6.239	Subversion: Create Remote Directory.....	6-142
6.240	Subversion: Create Subversion Repository.....	6-142
6.241	Subversion: Create/Edit Subversion Connection.....	6-143
6.242	Subversion: Delete Resources	6-143
6.243	Subversion: Edit Configuration File.....	6-144
6.244	Subversion: Export Files	6-144

6.245	Subversion: Export Subversion Connections.....	6-144
6.246	Subversion: Ignore.....	6-144
6.247	Subversion: Import Subversion Connections.....	6-145
6.248	Subversion: Import to Subversion.....	6-145
6.248.1	Destination.....	6-145
6.248.2	Source.....	6-145
6.248.3	Filters.....	6-145
6.248.4	Options.....	6-145
6.248.5	Summary.....	6-146
6.249	Subversion: Lock Resources.....	6-146
6.250	Subversion: Merge.....	6-146
6.251	Subversion: Pending Changes.....	6-147
6.252	Subversion: Properties.....	6-147
6.253	Subversion: Remove from Subversion.....	6-148
6.254	Subversion: Repository Browser.....	6-148
6.255	Subversion: Revert Local Changes.....	6-148
6.256	Subversion: Switch.....	6-148
6.257	Subversion: Unlock Resources.....	6-149
6.258	Subversion: Update Resources.....	6-149
6.259	Subversion: Update Working Copy.....	6-149
6.260	Subversion: Versioning Properties.....	6-150
6.261	Third-Party Database Objects.....	6-150
6.262	Unable to Connect.....	6-150
6.263	Unable to Open File.....	6-150
6.264	Uninstall Oracle REST Data Services.....	6-150
6.265	Unit Testing: Action Required.....	6-151
6.266	Unit Testing: Add Category.....	6-151
6.267	Unit Testing: Add Data Type.....	6-151
6.268	Unit Testing: Add Item to Library.....	6-151
6.269	Unit Testing: Add Test Implementation.....	6-152
6.270	Unit Testing: Add Test Suite.....	6-152
6.271	Unit Testing: Add Tests or Suites to a Suite.....	6-152
6.272	Unit Testing: Copy or Rename Unit Test.....	6-152
6.273	Unit Testing: Create Unit Test.....	6-152
6.273.1	Select Operation.....	6-153
6.273.2	Specify Test Name.....	6-153
6.273.3	Specify Startup.....	6-153
6.273.4	Specify Parameters.....	6-153
6.273.5	Specify Validations.....	6-153
6.273.6	Specify Teardown.....	6-154
6.273.7	Summary.....	6-154
6.274	Unit Testing: Manage Users.....	6-155
6.275	Unit Testing: Rename Test Implementation.....	6-155

6.276	Unit Testing: Result of Operation.....	6-155
6.277	Unit Testing: Synchronize Unit Test	6-155
6.278	Unplug Pluggable Database.....	6-156
6.279	Unsupported Database Version.....	6-156
6.280	Validate Geometry.....	6-156
6.281	View Whole Value	6-157
6.282	Windows	6-157
6.283	XMLType Column Properties.....	6-157

7 SQL Developer: TimesTen Support

7.1	Create a New TimesTen Connection	7-2
7.1.1	DSN	7-3
7.1.2	Easy Client/Server	7-4
7.1.3	Advanced.....	7-5
7.2	Set the Cache Administrator User Name and Password	7-5
7.3	Start and Stop the Cache Agent	7-6
7.4	Create a Cache Group	7-6
7.4.1	Defining a cache group and cache table attributes.....	7-7
7.4.2	Defining the root table.....	7-8
7.4.3	Defining the child tables.....	7-11
7.5	Drop a Cache Group.....	7-13
7.6	Load or Refresh a Cache Group.....	7-13
7.7	Flush a Cache Group	7-14
7.8	Unload a Cache Group.....	7-14
7.9	Alter the Automatic Refresh Attributes of a Cache Group	7-15
7.10	Specify and Display LOB Data Types.....	7-15
7.11	Specify the INLINE Attribute for Columns	7-15
7.12	Compute Table Size.....	7-16
7.13	Use a REF CURSOR as an OUTPUT Parameter.....	7-17
7.14	Create an Index on a TimesTen Table.....	7-17
7.15	Create a Primary Key Constraint	7-18
7.16	Change the Attributes of Primary Key	7-18
7.17	Create a Foreign Key Constraint	7-19
7.18	Drop a Constraint	7-20
7.19	Add an Aging Policy to a Table.....	7-20
7.20	Drop an Aging Policy from a Table	7-21
7.21	Change the Aging State for the Aging Policy of a Table.....	7-21
7.22	Change the Memory Usage Thresholds and LRU Aging Cycle	7-21
7.23	Change the Lifetime and Cycle for the Time-Based Aging Policy of a Table.....	7-21
7.24	Grant Object Privileges to a TimesTen Database Element	7-22
7.25	Revoke Object Privileges from a TimesTen Database Element	7-22
7.26	Load Data into a Table	7-23
7.27	Start and Stop the TimesTen Replication Agent	7-23

7.28	Update the Table and Column Statistics Used By the TimesTen Query Optimizer	7-23
7.29	Export the Table Statistics Used By the TimesTen Query Optimizer	7-25
7.30	Create a TimesTen Performance Metrics Snapshot	7-25
7.31	Drop a TimesTen Performance Metrics Snapshot	7-26
7.32	Edit the TimesTen Performance Metrics Configuration	7-27
7.33	Create a TimesTen Performance Metrics Report	7-27
7.33.1	Save a report.....	7-28
7.34	Capture Data for the TimesTen Index Advisor at the Connection Level.....	7-29
7.34.1	Capture data from a SQL worksheet.....	7-29
7.34.2	Capture data from a SQL workload script	7-31
7.35	Capture Data for the TimesTen Index Advisor at the Database Level.....	7-33
7.36	View the TimesTen Index Advisor Capture State	7-35
7.37	Load Data from an Oracle Database into a TimesTen Database	7-36
7.38	Load Data from an Oracle Database Table into an Existing TimesTen Table	7-39

8 SQL Developer Accessibility Information

8.1	About Oracle SQL Developer Accessibility	8-1
8.2	Using a Screen Reader and Java Access Bridge with Oracle SQL Developer.....	8-1
8.3	Oracle SQL Developer Features that Support Accessibility	8-1
8.3.1	Keyboard Access	8-2
8.3.2	Screen Reader Readability	8-2
8.3.3	Flexibility in Font and Color Choices.....	8-3
8.3.4	No Audio-only Feedback	8-3
8.3.5	No Dependency on Blinking Cursor and Animation	8-3
8.3.6	Screen Magnifier Usability.....	8-3
8.3.7	How to Change the Editor or Tabbed View of a File.....	8-4
8.3.8	How to Read Text in a Multiline Edit Field.....	8-4
8.3.9	How to Read the Line Number in the Source Editor	8-4
8.3.10	How to Access Exception Stack HTML Links and Generated Javadoc Links in the Log Window	8-4
8.4	Recommendations for Customizing Oracle SQL Developer.....	8-4
8.4.1	How to Customize the Accelerators Keys	8-4
8.4.2	How to Pass a Conflicting Accelerator Key to Oracle SQL Developer	8-4
8.4.3	How to Change the Look and Feel of SQL Developer.....	8-4
8.4.4	How to Customize the Fonts in Code Editors.....	8-5
8.4.5	How to Customize Syntax Highlighting.....	8-5
8.4.6	How to Display Line Numbers in Code Editors.....	8-5
8.4.7	How to Change the Timing for Completion Insight	8-5
8.4.8	How to Specify the Columns in the Debugger	8-5
8.5	Highly Visual Features of Oracle SQL Developer	8-5

Index

List of Figures

1-1	SQL Developer Main Window.....	1-4
2-1	V-model with a Database Migration.....	2-26
2-2	Main Window for a Database Migration.....	2-32
4-1	Unit Test Navigator.....	4-3

List of Tables

1-1	Default Locations for User-Related Information.....	1-122
-----	---	-------

Preface

This guide provides conceptual and usage information about Oracle SQL Developer, a graphical tool that enables you to browse, create, edit, and delete (drop) database objects; run SQL statements and scripts; edit and debug PL/SQL code; manipulate and export data; migrate third-party databases to Oracle; view metadata and data in MySQL and third-party databases; and view and create reports.

Audience

This guide is intended for those using the Oracle SQL Developer tool.

Documentation Accessibility

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

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Product Accessibility

If you are using a screen reader (such as JAWS), SQL Developer must be running on the same system as the screen reader.

For more information about using a screen reader and Java Access Bridge with SQL Developer, see [SQL Developer Accessibility Information](#) and *Oracle SQL Developer Installation Guide*.

Related Documents

For information about installing Oracle SQL Developer, see the *Oracle SQL Developer Installation Guide*.

To download free release notes, installation documentation, white papers, or other collateral, go to the Oracle Technology Network (OTN) at

<http://www.oracle.com/technetwork/>

The documentation section of the OTN site is at

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	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.

Third-Party License Information

Oracle SQL Developer contains third-party code. Oracle is required to provide the following notices. Note, however, that the Oracle program license that accompanied this product determines your right to use the Oracle program, including the third-party software, and the terms contained in the following notices do not change those rights.

Antlr 3.1.3

This program contains Antlr 3.1.3, which has the following associated attribution text:

ANTLR 3 License [The BSD License] Copyright (c) 2003-2007, Terence Parr All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. * Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Batik Version 1.7

This program contains Batik version 1.7, which has the following associated attribution text:

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>). This software contains code from the World Wide Web Consortium (W3C) for the Document Object Model API (DOM API) and SVG Document Type Definition (DTD). This software contains code from the International Organisation for Standardization for the definition of character entities used in the software's documentation.

2. Must include Apache License, Version 2.0 with any distribution of Apache Batik 1.7: Apache License Version 2.0, January 2004 <http://www.apache.org/licenses/> TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions. "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document. "Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License. "Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity. "You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License. "Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files. "Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types. "Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below). "Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof. "Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution." "Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by

combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed. 4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions: a. You must give any other recipients of the Work or Derivative Works a copy of this License; and b. You must cause any modified files to carry prominent notices stating that You changed the files; and c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License. 5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions. 6. Trademarks.

This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file. 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License. 8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages. 9. Accepting Warranty or Additional Liability. While redistributing the Work or

Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability. END OF TERMS AND CONDITIONS Appendix: How to Apply Apache License to Your Work To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives. Copyright [yyyy] [name of copyright owner] Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0> Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Bouncy Castle Crypto API 1.45

This program includes Bouncy Castle Crypto API 1.45, which has the following associated attribution text:

Copyright (c) 2000 - 2009 The Legion Of The Bouncy Castle (<http://www.bouncycastle.org/>) Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

JGoodies Looks 1.3.1

This program includes JGoodies Looks 1.3.1, which has the following associated attribution text:

The BSD License for the JGoodies Looks

Copyright © 2003 JGoodies Karsten Lentzsch. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of JGoodies Karsten Lentzsch nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

JXLayer 4.0

This program includes JXLayer 4.0, which has the following associated attribution text:

Copyright (c) 2006-2009, Alexander Potochkin All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of the JXLayer project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Log4j 1.2.9

This program contains Log4j 1.2.9, which has the following associated attribution text:

The Apache Software License, Version 1.1 * * Copyright (c) 2000 The Apache Software Foundation. All rights * reserved. * * Redistribution and use in source and binary forms, with or without * modification, are permitted provided that the following conditions * are met: * * 1. Redistributions of source code must retain the above copyright * notice, this list of conditions and the following disclaimer. * * 2. Redistributions in binary form must reproduce the above copyright * notice, this list of conditions and the following disclaimer in * the documentation and/or other materials provided with the * distribution. * * 3. The end-user documentation included with the redistribution, * if any, must include the following acknowledgment: * "This product includes software developed by the * Apache Software Foundation (Link1 /)." * Alternately, this acknowledgment may appear in the software itself, * if and wherever such third-party acknowledgments normally appear. * * 4. The names "Apache" and "Apache Software Foundation" must * not be used to endorse or promote products derived from this * software without prior written permission. For written * permission, please contact apache@apache.org. * * 5. Products derived from this software may not be called "Apache", * nor may "Apache" appear in their name, without prior written * permission of the Apache Software Foundation. * * THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED * WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES * OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE * DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR * ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, * SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT * LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF * USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND * ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, * OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT * OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF * SUCH DAMAGE. *

=====
* * This software consists of voluntary contributions made by many * individuals on behalf of the Apache Software Foundation. For more * information on the Apache Software Foundation, please see * <Link2 />. * * Portions of this software are based upon public domain software * originally written at the National Center for Supercomputing Applications, * University of Illinois, Urbana-Champaign. */

POI 2.5.1

This program contains POI 2.5.1, which has the following associated attribution text:

The following applies to all products licensed under the Apache 2.0 License: You may not use the identified files except in compliance with the Apache License, Version 2.0 (the "License.") You may obtain a copy of the License at <http://www.apache.org/>. A copy of the license is also reproduced below. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License. Apache License Version 2.0, January 2004 Link2 / TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION 1. Definitions. "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document. "Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License. "Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause

the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity. "You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License. "Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files. "Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types. "Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below). "Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof. "Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner.

For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution." "Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions: (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and (b) You must cause any modified files to carry prominent notices stating that You changed the files; and (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable

copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License. 5. Submission of Contributions.

Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions. 6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file. 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License. 8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages. 9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability. END OF TERMS AND CONDITIONS APPENDIX: How to apply the Apache License to your work. To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives. Copyright [yyyy] [name of copyright

owner] Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at [Link3](#) Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License. Oracle must give recipients a copy of this license (doc).

Regular Expression Package 1.3

This program includes Regular Expression Package, which has the following associated attribution text:

Jakarta Regexp Version 1.3, Copyright 1999-2002 The Apache Software Foundation. All rights reserved. Licensed under the Apache 1.1 License Agreement. The names "The Jakarta Project", "Jakarta-Regexp", and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org. This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

The Apache Software License, Version 1.1 Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment: "This product includes software developed by the Apache Software Foundation ([Link2](#) /)." Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear. 4. The names identified above with the specific software must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org. 5. Products derived from this software may not be called "Apache" nor may "Apache" appear in their names without prior written permission of the Apache Group. THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

SQL Developer Concepts and Usage

This topic covers the main Oracle SQL Developer concepts, and explains how to use the major SQL Developer features.

This topic includes the following topics:

- [About SQL Developer](#) (page 1-2)
- [Installing and Getting Started with SQL Developer](#) (page 1-2)
- [SQL Developer User Interface](#) (page 1-3)
- [Database Objects](#) (page 1-13)
- [Database Connections](#) (page 1-28)
- [Entering and Modifying Data](#) (page 1-34)
- [Running and Debugging Functions and Procedures](#) (page 1-36)
- [Using the SQL Worksheet](#) (page 1-41)
- [Using Snippets to Insert Code Fragments](#) (page 1-50)
- [Finding Database Objects](#) (page 1-52)
- [Using Recent Objects](#) (page 1-53)
- [Using Versioning](#) (page 1-53)
- [Using DBA Features in SQL Developer](#) (page 1-54)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)
- [Using the Cart](#) (page 1-68)
- [Spatial Support in SQL Developer](#) (page 1-71)
- [Change Manager Support in SQL Developer](#) (page 1-74)
- [Apache Hadoop Connectors Support in SQL Developer](#) (page 1-79)
- [SQL Developer Reports](#) (page 1-82)
- [SQL Developer Preferences](#) (page 1-92)
- [Location of User-Related Information](#) (page 1-122)
- [Java Development Kit \(JDK\) Location](#) (page 1-123)
- [Command-Line Interface for SQL Developer](#) (page 1-123)

- [Data Modeler in SQL Developer](#) (page 1-124)
- [Oracle TimesTen In-Memory Database Support](#) (page 1-125)
- [Using the Help](#) (page 1-125)
- [SQL Developer Usage Suggestions](#) (page 1-126)
- [For More Information About SQL Developer](#) (page 1-129)

Related Topics:

[SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

[SQL Developer: Oracle REST Data Services Support](#) (page 3-1)

[SQL Developer: Unit Testing](#) (page 4-1)

[SQL Developer Tutorial: Creating Objects for a Small Database](#) (page 5-1)

1.1 About SQL Developer

Oracle SQL Developer is a graphical version of SQL*Plus that gives database developers a convenient way to perform basic tasks. You can browse, create, edit, and delete (drop); run SQL statements and scripts; edit and debug PL/SQL code; manipulate and export (unload) data; and view and create reports.

You can connect to any target Oracle Database schema using standard Oracle Database authentication. Once connected, you can perform operations on objects in the database.

You can connect to schemas for MySQL and selected third-party (non-Oracle) databases, such as Microsoft SQL Server, Sybase Adaptive Server, and IBM DB2, and view metadata and data in these databases; and you can migrate these databases to Oracle Database.

1.2 Installing and Getting Started with SQL Developer

To install and start SQL Developer, you simply download a ZIP file and unzip it into a desired parent directory or folder, and then type a command or double-click a file name. You should read the *Oracle SQL Developer Installation Guide* before you perform the installation. After you have read the installation guide, the basic steps are:

1. Unzip the SQL Developer kit into a folder (directory) of your choice, which will be referred to as `<sqldeveloper_install>`. Unzipping the SQL Developer kit causes a folder named `sqldeveloper` to be created under the `<sqldeveloper_install>` folder.

For example, if you unzip the kit into `C:\`, the folder `C:\sqldeveloper` is created, along with several subfolders under it.

If Oracle Database (Release 11 or later) is also installed, a version of SQL Developer is also included and is accessible through the menu system under Oracle. This version of SQL Developer is separate from any SQL Developer kit that you download and unzip on your own, so do not confuse the two, and do not unzip a kit over the SQL Developer files that are included with Oracle Database. Suggestion: Create a shortcut for the SQL Developer executable file that you install, and always use it to start SQL Developer.

2. To start SQL Developer, go to the `sqldeveloper` directory under the `<sqldeveloper_install>` directory, and do one of the following:

On Linux and Mac OS X systems, run `sh sqldeveloper.sh`.

On Windows systems, double-click `sqldeveloper.exe`.

If you are asked to enter the full pathname for the JDK, click **Browse** and find it. For example, on a Windows system the path might have a name similar to `C:\Program Files\Java\jdk1.7.0_51`. (If you cannot start SQL Developer, it could be due to an error in specifying or configuring the JDK.)

3. If you want to become familiar with SQL Developer concepts before using the interface, read the rest of this chapter before proceeding to the next step.
4. Create at least one database connection (or import some previously exported connections), so that you can view and work with database objects, use the SQL Worksheet, and use other features.

To create a new database connection, right-click the **Connections** node in the Connections navigator, select **New Connection**, and complete the required entries in the Create/Edit/Select Database Connection dialog box. (You may also be able to generate connections automatically by right-clicking the Connections node and selecting **Create Local Connections**.)

5. Do the short tutorial in [SQL Developer Tutorial: Creating Objects for a Small Database](#) (page 5-1), or work with your existing objects.

Related Topics:

[Java Development Kit \(JDK\) Location](#) (page 1-123)

[Database Connections](#) (page 1-28)

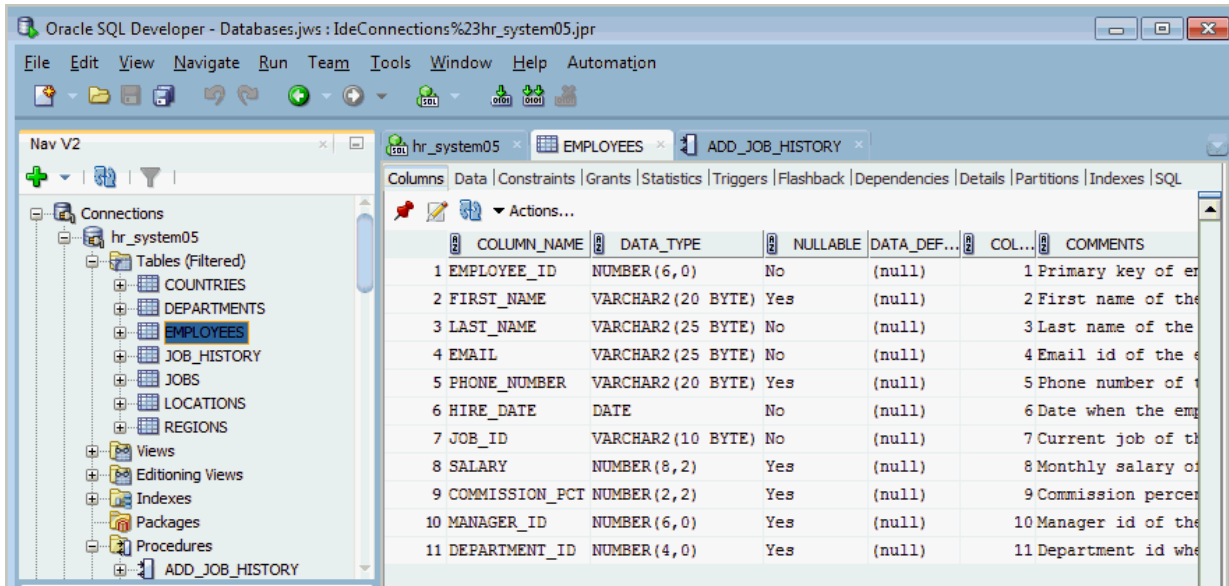
[Create/Edit/Select Database Connection](#) (page 6-11)

1.3 SQL Developer User Interface

The SQL Developer window generally uses the left side for navigation to find and select objects, and the right side to display information about selected objects.


[Figure 1-1](#) (page 1-4) shows the main window.

Figure 1-1 SQL Developer Main Window

**Note:**

This text explains the default interface. However, you can customize many aspects of the appearance and behavior of SQL Developer by setting preferences. If you ever need to restore the default interface, see [Restoring the Original "Look and Feel"](#) (page 1-12).

The menus at the top contain standard entries, plus entries for features specific to SQL Developer as shown in the following figure.



You can use **shortcut keys** to access menus and menu items: for example Alt+F for the File menu and Alt+E for the Edit menu; or Alt+H, then Alt+S for Help, then Search. You can also display the File menu by pressing the F10 key (except in the SQL Worksheet, where F10 is the shortcut for Explain Plan).

To close a window that has focus (such as the SQL Developer main window, a wizard or dialog box, or the Help Center) and any of its dependent windows, you can press Alt+F4.

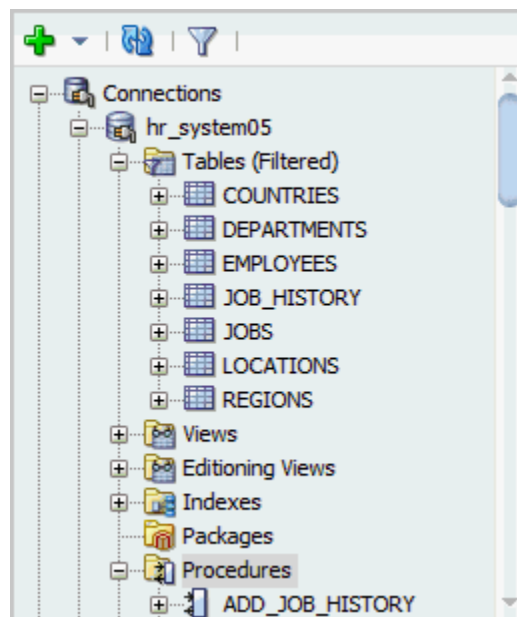
The main toolbar (under the menus) contains Icons to perform various actions, which by default include the following:

- **New** creates a database object.
- **Open** opens a file.
- **Save** saves any changes to the currently selected object.
- **Save All** saves any changes to all open objects.

- **Back** moves to the pane that you most recently visited. (Or use the drop-down arrow to specify a tab view.)
- **Forward** moves to the pane after the current one in the list of visited panes. (Or use the drop-down arrow to specify a tab view.)
- **Open SQL Worksheet** opens the SQL Worksheet. If you do not use the drop-down arrow to specify the database connection to use, you are asked to select a connection.

However, you can add and remove main toolbar icons by clicking **View, Show Toolbars, Customize Main Toolbar**, then choosing desired options. To undo any main toolbar customization, select **Reset Toolbars**.

The left side of the SQL Developer window has panes for the Connections and Reports navigators (any other navigators that you might open), icons for performing actions, and a hierarchical tree display for the currently selected navigator, as shown in the following figure.



The **Connections navigator** lists database connections that have been created. To create a new database connection, import an XML file with connection definitions, or export or edit current connections, right-click the Connections node and select the appropriate menu item.

The **Files navigator** (marked by a folder icon; not shown in the preceding figure) displays your local file system using a standard hierarchy of folders and files. You can double-click or drag and drop files to open them, and you can edit and save the files. For example, if you open a .sql file, it is displayed in a SQL Worksheet window. The Files navigator is especially useful if you are using versioning with SQL Developer.

The **Reports navigator** (not shown in the preceding figure) lists informative reports provided by SQL Developer, such as a list of tables without primary keys for each database connection, as well as any user-defined reports.

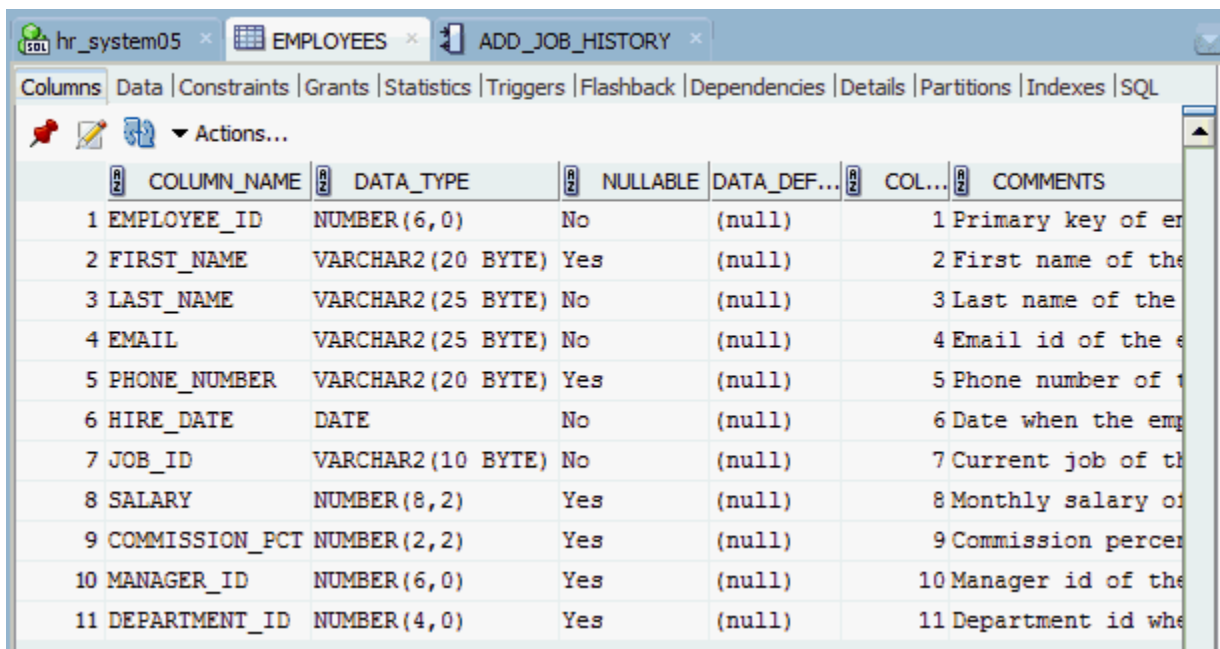
Icons above the Connections node and metadata tree perform the following actions on the currently selected object:

- **Refresh** queries the database for the current details about the selected object (for example, a connection or just a table).
- **Apply Filter** restricts the display of objects using a filter that you specify. For example, you can right-click the Tables node and specify a filter of EM% to see only tables that start with EM and to have the Tables node label be changed to *Tables (EM%)*. To remove the effects of applying a filter, right-click the node and select **Clear Filter**.

Note that for tables, the initial default node label is *Tables (Filtered)*, to reflect the exclusion of tables that are in the Recycle Bin. To have such tables included in the display, clear the filter.

The metadata tree in the Connections pane displays all the objects (categorized by object type) accessible to the defined connections. To select an object, expand the appropriate tree node or nodes, then click the object.

The right side of the SQL Developer window has tabs and panes for objects that you select or open, as shown in the following figure, which displays information about a table named EMPLOYEES. (If you hold the mouse pointer over the tab label -- EMPLOYEES in this figure -- a tooltip displays the object's owner and the database connection.)



The screenshot shows the SQL Developer interface with the 'EMPLOYEES' table selected. The 'Columns' pane displays the following table structure:

	COLUMN_NAME	DATA_TYPE	NULLABLE	DATA_DEF...	COL...	COMMENTS
1	EMPLOYEE_ID	NUMBER (6,0)	No	(null)	1	Primary key of er
2	FIRST_NAME	VARCHAR2 (20 BYTE)	Yes	(null)	2	First name of the
3	LAST_NAME	VARCHAR2 (25 BYTE)	No	(null)	3	Last name of the
4	EMAIL	VARCHAR2 (25 BYTE)	No	(null)	4	Email id of the e
5	PHONE_NUMBER	VARCHAR2 (20 BYTE)	Yes	(null)	5	Phone number of t
6	HIRE_DATE	DATE	No	(null)	6	Date when the emp
7	JOB_ID	VARCHAR2 (10 BYTE)	No	(null)	7	Current job of th
8	SALARY	NUMBER (8,2)	Yes	(null)	8	Monthly salary of
9	COMMISSION_PCT	NUMBER (2,2)	Yes	(null)	9	Commission perce
10	MANAGER_ID	NUMBER (6,0)	Yes	(null)	10	Manager id of the
11	DEPARTMENT_ID	NUMBER (4,0)	Yes	(null)	11	Department id whe

For objects other than subprograms, icons provide the following options:

- **Freeze View** (the **pin**) keeps that object's tab and information in the window when you click another object in the Connections navigator; a separate tab and display are created for that other object. If you click the pin again, the object's display is available for reuse.
- **Edit** displays a dialog box for editing the object.
- **Refresh** updates the display by querying the database for the latest information.
- **Actions** displays a menu with actions appropriate for the object. The actions are the same as when you right-click an object of that type in the Connections navigator, except the Actions menu does not include Edit.

To switch among objects, click the desired tabs; to close a tab, click the X in the tab. If you make changes to an object and click the X, you are asked if you want to save the changes.

For tables and views, this information is grouped under tabs, which are labeled near the top. For example, for tables the tabs are Columns, Data (for seeing and modifying the data itself), Indexes, Constraints, and so on; and you can click a column heading under a tab to sort the grid rows by the values in that column. For most objects, the tabs include SQL, which displays the SQL statement for creating the object.

You can export (unload) data from a detail pane or from the results of a SQL Worksheet operation or a report by using the context menu and selecting **Export**.

The **Messages - Log** area is used for feedback information as appropriate (for example, results of an action, or error or warning messages). If this area is not already visible, you can display it by clicking View and then Log.

The **Compiler - Log** area is used for any messages displayed as a result of a Compile or Compile for Debug operation.

Related Topics:

[Menus for SQL Developer](#) (page 1-7)

[SQL Developer Preferences](#) (page 1-92)

[Using the SQL Worksheet](#) (page 1-41)

[SQL Developer Reports](#) (page 1-82)

1.3.1 Menus for SQL Developer

This topic explains menu items that are specific to SQL Developer.

File menu

Data Modeler: Lets you open, close, import, export, and print Data Modeler designs, and use Data Modeler reports (see [Data Modeler in SQL Developer](#) (page 1-124)).

Edit menu

Extended Paste: Displays the [Paste](#) (page 6-117) dialog box, in which you select a clipboard item (from potentially many) to be pasted into the current location.

Duplicate Selection: When you have selected text while editing a function or procedure, creates a copy of the selected text at the current location.

Wrap Selection: When you have selected text while editing a function or procedure, wraps the selected text.

View menu

Contains options that affect what is displayed in the SQL Developer interface.

Breakpoints: Displays the Breakpoints pane, which shows breakpoints, both system-defined and user-defined (see [Running and Debugging Functions and Procedures](#) (page 1-36)).

Cart: Displays the Cart window (see [Using the Cart](#) (page 1-68)).

Change Management: Displays the Change Management window (see [Change Manager Support in SQL Developer](#) (page 1-74)).

Components: Displays the [Component Palette: Configure Component Palette](#) (page 6-8) dialog box.

Connections: Displays the Connections navigator.

DBA: Displays the DBA navigator (see [Using DBA Features in SQL Developer](#) (page 1-54)).

Data Miner: Lets you display the Data Miner Navigator, Workflow Jobs, Workflow Property Inspector, and Component Palette. (For information about Oracle Data Miner, click **Help**, then **Data Mining**.)

Data Modeler: Lets you display the Browser and Thumbnail Diagram panes of the [Data Modeler in SQL Developer](#) (page 1-124).

DBMS Output: Displays the output of DBMS_OUTPUT package statements (see [DBMS Output Pane](#) (page 1-46)).

Debugger: Displays panes related to debugging (see [Running and Debugging Functions and Procedures](#) (page 1-36)).

Extension Diagnostics: Displays the Extensions - Log pane.

Files: Displays the Files navigator, which is marked by a folder icon. You can use the Files navigator to browse, open, edit, and save files that are accessible from the local system.

Find DB Object: Displays the Find Database Object pane (see [Finding Database Objects](#) (page 1-52)).

Log: Displays a pane with these tabs:

- **Messages - Log**, which can contain errors, warnings, and informational messages.
- **Logging Page**, with detailed information about individual log entries. You can right-click in the page to *change the logging level*, for example, from Severe to warning. (Going down the list of logging levels, each provides more information than the preceding one and includes all information provided by all preceding levels, with All providing the most information.)

Map View: Displays the Map View pane (see [Map Visualization of Spatial Data](#) (page 1-72)).

Migration Projects: Lets you display any captured models and converted models (see [SQL Developer User Interface for Migration](#) (page 2-32)).

OWA Output: Displays Oracle Web Agent (MOD_PLSQL) output (see [OWA Output Pane](#) (page 1-47)).

Processes: Displays the Processes pane.

Properties: Displays the Property Inspector.

REST Data Services: Lets you display the REST Development pane (see [REST Development](#) (page 3-12)) or the ORDS Administration pane (see [Oracle REST Data Services Administration](#) (page 3-1)).

Recent Objects: Displays a pane with names of recently opened objects. You can double-click a name in the list to go to its editing window (see [Using Recent Objects](#) (page 1-53)).

Reports: Displays the Reports navigator (see [SQL Developer Reports](#) (page 1-82)).

SQL History: Displays information about SQL statements that you have executed. You can select statements and append them to or overwrite statements on the worksheet (see [SQL History](#) (page 1-47)).

Snippets: Displays snippets (see [Using Snippets to Insert Code Fragments](#) (page 1-50)).

Task Progress: Displays the Task Progress pane.

Unit Test: Displays the Unit Test navigator (see [SQL Developer: Unit Testing](#) (page 4-1)).

Show Status Bar: Controls the display of the status bar at the bottom of the SQL Developer window.

Show Toolbars: Controls the display of the following toolbars: Main toolbar (under the SQL Developer menus), the Connections Navigator toolbar, the DBA navigator toolbar, the Data Miner Workflow Property Inspector toolbar, and (if a package or subprogram is open) the Code Editor toolbar.

OLAP: Lets you display the OLAP Templates and Snippets pane.

Navigate menu

Contains options for navigating to panes and in the execution of subprograms.

Back: Moves to the pane that you most recently visited.

Forward: Moves to the pane after the current one in the list of visited panes.

Toggle Bookmark: If you are editing a function or procedure, creates or removes a bookmark (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Remove Bookmarks from File: Removes bookmarks from the currently active editing window for a function or procedure (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Remove All Bookmarks: Removes bookmarks from open editing windows for functions and procedures (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Go to Bookmark: Displays a dialog box so that you can go to a specified bookmark (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Go to Next Bookmark: Goes to the next bookmark in the currently active editing window for a function or procedure (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Go to Previous Bookmark: Goes to the previous bookmark in the currently active editing window for a function or procedure (see [Using Bookmarks When Editing Functions and Procedures](#) (page 1-39)).

Go to Line: Goes to the specified line number and highlights the line in the editing window for the selected function or procedure. (To display line numbers, enable **Show Line Numbers** under the [Code Editor](#) (page 1-94): **Line Gutter** preferences.)

Go to Last Edit: Goes to the last line that was edited in the editing window for a function or procedure.

Go to Recent Files: Displays the [Recent Files](#) (page 6-122) dialog box, in which you can specify a function or procedure to go to.

Run menu

Contains options relevant when a function or procedure is selected or when it is open for debugging.

Run [name]: Starts execution of the specified function or procedure.

Debug [name]: Starts execution of the specified function or procedure in debug mode.

The remaining items on the Debug menu match commands on the debugging toolbar, which is described in [Running and Debugging Functions and Procedures](#) (page 1-36).

Team menu

Contains options related to support for the Subversion version management and source control system, and for any other such systems (such as CVS) that you have added as extensions to SQL Developer through the "check for updates" feature. See [Using Versioning](#) (page 1-53) for more information.

The commands on the Team menu depend on which version management and source control systems are available for use with SQL Developer.

Tools menu

Invokes SQL Developer tools.

Data Modeler: Starts the [Data Modeler in SQL Developer](#) (page 1-124) if it not already active; otherwise, contains the commands **About Data Modeler**, **Design Rules**, and **General Options** (user preferences).

Database Copy: Enables you to copy objects, schemas, or a tablespace from one database to another (see the [Database Copy](#) (page 6-73) wizard).

Database Diff: Enables you to compare two schemas to find differences between objects of the same type and name (for example, tables named CUSTOMERS) in two different schemas, and optionally to update the objects in the destination schema to reflect differences in the source schema (see the [Database Differences](#) (page 6-75) interface).

Database Export: Enables you to export (unload) some or all objects of one or more object types for a database connection to a file containing SQL statements to create these objects and optionally to export table data (see the [Database Export \(Unload Database Objects and Data\)](#) (page 6-77) interface).

Migration: Displays the [Migration Submenu](#) (page 2-33), which contains commands related to migrating third-party databases to Oracle.

Real Time SQL Monitor (requires the Oracle Tuning Pack): Displays information about any query currently executing and queries that are done executing for a selected connection. To see detailed information about a query, right-click its row and select Show SQL Details. The information is especially useful for real-time monitoring of long-running SQL statements. Cursor statistics (such as CPU times and IO times) and execution plan statistics (such as number of output rows, memory, and temporary space used) are updated close to real-time during statement execution. (Internally, this feature calls the DBMS_SQLTUNE.REPORT_SQL_MONITOR subprogram.)

Monitor Sessions: Displays the status of one or more sessions, using information from the V\$RSRC_SESSION_INFO view, which shows how the session has been affected by the Oracle Database Resource Manager.

SQL Worksheet: Displays a worksheet in which you can enter and execute SQL and PL/SQL statements using a specified connection (see [Using the SQL Worksheet](#) (page 1-41)).

Unit Test: Displays the [Unit Test Submenu](#) (page 4-3), which contains commands related to unit testing.

Data Miner: Enables you to show the Data Miner navigator and drop the Data Miner repository. (For information about Data Miner, click **Help**, then **Data Mining**).

OLAP: Enables you to create a view that pivots an Oracle OLAP fact table so that the measures identify rows instead of columns. (For information about Oracle OLAP, click **Help** and check for related topics in the Table of Contents.)

REST Data Services: Lets you manage connections for REST administration and development, and install and uninstall Oracle REST Data Services.

External Tools: Displays the [External Tools](#) (page 6-87) dialog box, with information about user-defined external tools that are integrated with the SQL Developer interface. From this dialog box can add external tools (see [Create/Edit External Tool](#) (page 6-88)). The Tools menu also contains items for any user-defined external tools.

Features: Displays the [Manage Features and Updates](#) (page 6-112) dialog box.

Preferences: Enables you to customize the behavior of SQL Developer (see [SQL Developer Preferences](#) (page 1-92)).

Window menu

Contains options relating to the appearance and behavior of the user interface.

Team: Lets you display the Versions and Pending Changes.

Switch Window Layout: In relevant contexts, lets you switch between editing and debugging mode.

Configure Window: Lets you perform operations on the selected window or window group, such as maximizing, floating, minimizing, and docking.

Reset Windows to Factory Settings: Restores the SQL Developer interface to the way it was after SQL Developer was installed; however, this does not undo any main toolbar customizations. (See also [Restoring the Original "Look and Feel"](#) (page 1-12), and Reset Toolbars in the [Customize Toolbars](#) (page 6-70) dialog box.)

Assign File Accelerator: Lets you create [SQL Worksheet Bookmarks](#) (page 1-126).

Close Window: Closes the currently selected window in the SQL Worksheet.

Close All Documents: Closes all open windows in the SQL Worksheet.

Close Other Documents: Closes all open windows in the SQL Worksheet except for the currently selected window.

Documents: Displays a Documents dialog box for performing operations on open windows.

Help menu

Displays help about SQL Developer and enables you to check for SQL Developer updates.

Search: Displays the Help Center window.

Table of Contents: Displays the Help Center window. In this window you can click these icons:

- **Keep on Top:** Toggles whether to keep the Help Center window on top of the SQL Developer window.
- **Navigators:** Lets you select a help navigator.

- **Print:** Prints the topic.
- **Change Font Size:** Lets you increase or decrease the font size for the display of the current help topic.
- **Add to Favorites:** Adds the topic to the Favorites list.
- **Find:** Lets you search for a string in the current help topic.

Start Page: Displays a page with links for options for learning about SQL Developer. It contains the tabs Get Started and Community (the latter including the SQL Developer discussion forum).

Data Mining: Displays an introductory help topic about Oracle Data Mining and its graphical user interface, Data Miner, which is integrated into SQL Developer.

Check for Updates: Checks for any updates to the selected optional SQL Developer extensions, as well as any mandatory SQL Developer extensions. The available updates may include the JTDS JDBC Driver for Microsoft SQL Server and the MySQL JDBC Driver, which enable you to create connections to databases other than Oracle Database. (If the system you are using is behind a firewall, see the SQL Developer user preferences for [Web Browser and Proxy](#) (page 1-121).)

About: Displays version-related information about SQL Developer and its components.

1.3.2 Restoring the Original "Look and Feel"

If you have made changes to the SQL Developer user interface ("look and feel"), such as accidentally repositioning navigators and panes, you can restore the window to the way it was after SQL Developer was installed by clicking **Window**, then **Reset Windows to Factory Settings**. (However, if you have customized the main toolbar and want to restore it, you must either also select **Reset Toolbars** in the Customize Toolbars dialog box or restore the entire user interface.)

To restore the entire user interface (for example, including undoing all changes to user preferences), follow these steps.

1. If you are running SQL Developer, exit.
2. Create a backup copy of the folder or directory where your SQL Developer user information is stored, in case you want to restore any old user-defined reports, snippets, code templates, or SQL history. The default location is a build-specific directory or folder under the following:
 - Windows: C:\Documents and Settings*<user-name>*\Application Data\SQL Developer
 - Linux or Mac OS X: ~/.sqldeveloper

If you have specified a nondefault location for your SQL Developer user information, create the backup copy of that folder or directory instead.

(If you do not want to use any old information or settings, you can skip creating a backup copy.)

3. Delete the original (not the backup) folder or directory where your user information is stored (explained in step 2).
4. Start SQL Developer.

This creates a folder or directory where your user information is stored (explained in step 2), which has the same content as when SQL Developer was installed.

If you have made changes to the SQL Developer shortcut key (accelerator key) mappings, you can restore the mappings to the defaults for your system by clicking **Tools**, then **Preferences**, then **Shortcut Keys**, then **More Actions**, then **Load Keyboard Scheme**, and then selecting **Default**.

Related Topics:

[Customize Toolbars](#) (page 6-70)

[Location of User-Related Information](#) (page 1-122)

1.4 Database Objects

You can create, edit, delete (drop), and add to the Cart most types of objects in an Oracle database by using the context menu (right-click, or Shift+F10) in the Connections navigator or by clicking the **Actions** button in the detail pane display. For some objects, you can do other operations, as appropriate for the object type.

For some object types the context menu includes **Open**, which generates a tabular overview display of information about objects of that type. For example, selecting **Open** for the Tables node in the Connections navigator displays a list of all tables for the connection, and for each table the table name, the number of rows, the table status, the number of columns, the comment (if any), and other information.

For some object types the context menu includes **Use as Template**, which displays the appropriate Create dialog box with information for the selected object filled in. You should then change the object name and any other aspects of the definition, as needed, to create a new object of that type.

Note:

The actions available from context menus and Actions buttons depend on the Oracle Database release number for the specified database connection. If an action mentioned in the text is not available with a connection, it may be that the feature was not available in that release of Oracle Database.

The dialog boxes for creating and editing many objects contain a tab or pane named **DDL** or **SQL**, where you can see the statement that SQL Developer will use to perform the actions you specify (CREATE to create an object, ALTER to apply changes to an existing object).

You can search for specific objects associated with an Oracle database connection by clicking the Search icon.

If you have connected to a MySQL database or a third-party (non-Oracle) database, such as Microsoft SQL Server, Sybase Adaptive Server, or IBM DB2, you can view their objects using the Connections navigator. (For information about connecting to MySQL and third-party databases, see the SQL Developer user preferences for Database: Third Party JDBC Drivers.)

Related Topics:

[Finding Database Objects](#) (page 1-52)

1.4.1 Applications (Application Express 3.0.1 and Later)

Effective with Oracle Application Express 3.0.1, if you use SQL Developer to connect to a schema that owns any Application Express applications, the Connections navigator has an **Application Express** node. You can click an application name to display tabs (Application, Pages, LOVs, Lists, Templates, Breadcrumbs, and so on) with information about the application.

You can perform the following operations on an Application Express application by right-clicking the application name in the Connections navigator and selecting an item from the menu:

- **Import Application:** Imports an application from a specified file and installs the application.
- **Deploy Application:** Deploys an application into a specified target schema.
- **Drop:** Deletes the application.
- **Modify Application:** Enables you to change the alias, name (Rename), status, global notification, and proxy server for the application.
- **Quick DDL:** Saves the DDL statements to create the application (or the selected component) to a file, a .zip file, a worksheet, or the system clipboard.
- **Refactor (in bulk):** Collects all anonymous blocks, refactors them into PL/SQL procedures, and places them in a package. The output of a refactor in bulk operation is a PL/SQL script, which you can review and save, and which you can execute to create the package.

The following operations are available only by right-clicking the Application Express node in the Connections navigator and selecting an item from the menu:

- **Start EPG:** Starts the embedded PL/SQL gateway for Application Express. Displays a dialog box for executing the following statements: `BEGIN DBMS_EPG.map_dad('APEX', '/apex/*'); end;`
- **Stop EPG:** Stops the embedded PL/SQL gateway for Application Express. Displays a dialog box for executing the following statements: `BEGIN DBMS_EPG.unmap_dad('APEX'); end;`

1.4.2 Cache Groups (Oracle TimesTen In-Memory Database)

A cache group describes a collection of in-memory database tables that map to all or a subset of the tables in an Oracle database. A cache group can consist of all or a subset of the rows and columns in these tables. Multiple cache groups can be used to cache different sets of related tables in the Oracle database.

1.4.3 Chains

A chain is an Oracle Scheduler object that enables you to implement dependency scheduling, in which jobs are started depending on the outcomes of one or more previous jobs. A chain consists of multiple steps that are combined using dependency rules. The dependency rules define the conditions that can be used to start or stop a step or the chain itself. Conditions can include the success, failure, or completion-codes or exit-codes of previous steps. Logical expressions, such as AND/OR, can be

used in the conditions. In a sense, a chain resembles a decision tree, with many possible paths for selecting which tasks run and when.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.4 Credentials

A credential is an Oracle Scheduler object that is a user name and password pair stored in a dedicated database object. A job uses a credential to authenticate itself with a database instance or the operating system so that it can run.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.5 Database Destinations

A database destination is an Oracle Scheduler object that defines a location for running a job. There are two types of destinations: an **external destination** specifies a remote host name and IP address for running a remote external job; a **database destination** specifies a remote database instance for running a remote database job.

If you specify a destination when you create a job, the job runs on that destination. If you do not specify a destination, the job runs locally, on the system on which it is created.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.6 Database Links (Public and Private)

A database link is a database object in one database that enables you to access objects on another database. The other database need not be an Oracle Database system; however, to access non-Oracle systems you must use Oracle Heterogeneous Services. After you have created a database link, you can use it to refer to tables and views in the other database. The Connections navigator has a **Database Links** node for all database links (public and private) owned by the user associated with the specified connection, and a **Public Database Links** node for all public database links on the database associated with the connection.

You can perform the following operations on a database link by right-clicking the database link name in the Connections navigator and selecting an item from the menu:

- **Test Database Link:** Validates the database link.
- **Drop:** Deletes the database link.

Related Topics:

[Create/Edit Database Link](#) (page 6-28)

1.4.7 Destination Groups

A destination group is an Oracle Scheduler object whose members are a list of Scheduler destination objects. In a **database destination group**, the members are

database destinations, for running remote database jobs. In an **external destination group**, members are external destinations, for running remote external jobs.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.8 Directories

A directory object specifies an alias for a directory (called a folder on Windows systems) on the server file system where external binary file LOBs (BFILEs) and external table data are located. To create a directory (that is, a directory object), you can use SQL Developer or the SQL statement `CREATE DIRECTORY`.

You can use directory names when referring to BFILEs in your PL/SQL code and OCI calls, rather than hard coding the operating system path name, for management flexibility. All directories are created in a single namespace and are not owned by an individual schema. You can secure access to the BFILEs stored within the directory structure by granting object privileges on the directories to specific users.

1.4.9 Editions

Edition-based redefinition, introduced in Oracle Database 11g Release 2 (11.2), 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 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.

To specify the current edition, right-click the edition name and select **Set Current Edition**. To create an edition under an existing edition, right-click the edition name and select **Create Edition**. To delete an edition (and optionally all editions under it), right-click the edition name and select **Drop Edition**.

Related Topics:

[Oracle Database Development Guide](#)

1.4.10 File Watchers

A file watcher is an Oracle Scheduler object that defines the location, name, and other properties of a file whose arrival on a system causes the Scheduler to start a job. You create a file watcher and then create a number of event-based jobs or event schedules that reference the file watcher. When the file watcher detects the arrival of the designated file, it raises a file arrival event. The job started by the file arrival event can retrieve the event message to learn about the newly arrived file.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.11 Functions

A function is a type of PL/SQL subprogram, which is a programming object that can be stored and executed in the database server, and called from other programming objects or applications. (Functions return a value; procedures do not return a value.)

You can perform the following operations on a function by right-clicking the function name in the Connections navigator and selecting an item from the menu:

- **Edit:** Displays the function text so that you can view and edit it.
- **Run:** Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the function in normal (not debug) mode.
- **Compile:** Performs a PL/SQL compilation of the function.
- **Compile for Debug:** Performs a PL/SQL compilation of the procedure, with PL/SQL library units compiled for debugging.
- **Profile** (for an Oracle Database Release 11.1 or later connection): Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the function and collects execution statistics.
- **Drop:** Deletes the function.
- **Grant:** Enables you to grant available privileges on the function to selected users.
- **Revoke:** Enables you to revoke available privileges on the function from selected users.
- **Format:** Reformats the text of the function definition.
- **Create Unit Test:** Creates a unit test for the function.
- **Quick DDL:** Enables you to export the DDL statement for creating the function to a file, a SQL Worksheet, or the clipboard.

Related Topics:

[Run/Debug/Profile PL/SQL](#) (page 6-132)

[SQL Developer: Unit Testing](#) (page 4-1)

1.4.12 Indexes

An index is a database object that contains an entry for each value that appears in the indexed column(s) of the table or cluster and provides direct, fast access to rows. Indexes are automatically created on primary key columns; however, you must create indexes on other columns to gain the benefits of indexing.

You can perform the following operations on an index by right-clicking the index name in the Connections navigator and selecting an item from the menu:

- **Drop:** Deletes the index.
- **Rebuild:** Recreates the index or one of its partitions or subpartitions. If the index is unusable, a successful rebuild operation makes the index usable. For a function-based index, rebuilding also enables the index; however, if the function on which the index is based does not exist, the rebuild operation fails.
- **Rename:** Changes the name of the index.
- **Make Unusable:** Prevents the index from being used by Oracle in executing queries. An unusable index must be rebuilt, or dropped and re-created, before it can be used again.

- **Coalesce:** Merges the contents of index blocks, where possible, to free blocks for reuse.
- **Compute Statistics:** For a function-based index, collects statistics on both the index and its base table using the DBMS_STATS package. Such statistics will enable Oracle Database to correctly decide when to use the index.
- **Rebuild Index Partition:** Rebuilds a specified index partition.
- **Drop Index Partition:** Drops (deletes) a specified index partition.
- **Quick DDL:** Saves the DDL statement to create the index to a file, a SQL Worksheet, or the system clipboard.

Related Topics:

[Create/Edit Index](#) (page 6-28)

1.4.13 Java Sources

Java sources can be created and managed in the database. You can create a Java source object by right-clicking the Java node in the Connections navigator, selecting **Load Java**, and specifying the Java source name and either entering the source code or loading a Java source, class, or resource from a file (BFILE). (A CREATE OR REPLACE AND RESOLVE JAVA SOURCE statement is executed using the information you specify.)

Related Topics:

Oracle Database Java Developer's Guide

1.4.14 Jobs

A job is an Oracle Scheduler object that is a collection of metadata that describes a user-defined task. It defines what needs to be executed (the action), when (the one-time or recurring schedule or a triggering event), where (the destinations), and with what credentials. A job has an owner, which is the schema in which it is created.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.15 Job Classes

A job class is an Oracle Scheduler object that enables the Scheduler administrator to group jobs for logical purposes, such as to assign the same set of attribute values to member jobs, to set service affinity for member jobs, to set resource allocation for member jobs, or to group jobs for prioritization.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.16 Operators

An operator is a type of PL/SQL subprogram, which is a programming object that can be stored and executed in the database server, and called from other programming objects or applications. Operators return a value.

You can perform the following operations on an operator by right-clicking the name in the Connections navigator and selecting an item from the menu:

- **Edit:** Displays the operator text so that you can view and edit it.
- **Run:** Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the operator in normal (not debug) mode.
- **Compile:** Performs a PL/SQL compilation of the operator.
- **Compile for Debug:** Performs a PL/SQL compilation of the operator, with PL/SQL library units compiled for debugging.
- **Profile** (for an Oracle Database Release 11.1 or later connection): Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the operator and collects execution statistics.

Related Topics:

[Run/Debug/Profile PL/SQL](#) (page 6-132)

1.4.17 Materialized Views

A materialized view is a database object that contains the results of a query. The FROM clause of the query can name tables, views, and other materialized views. Collectively these objects are called master tables (a replication term) or detail tables (a data warehousing term). This reference uses "master tables" for consistency. The databases containing the master tables are called the master databases.

Related Topics:

[Create/Edit View](#) (page 6-62)

[View Properties or Materialized View Properties pane](#) (page 6-63)

1.4.18 Materialized View Logs

A materialized view log is a table associated with the master table of a materialized view. When DML changes are made to master table data, Oracle Database stores rows describing those changes in the materialized view log and then uses the materialized view log to refresh materialized views based on the master table. This process is called incremental or fast refresh. Without a materialized view log, Oracle Database must reexecute the materialized view query to refresh the materialized view. This process is called a complete refresh. Usually, a fast refresh takes less time than a complete refresh.

1.4.19 Multitenant Container Database (CDB)

A multitenant container database (CDB) is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and nonschema objects that appears to an Oracle Net client as a non-CDB.

Related Topics:

Overview of Pluggable Databases

Architecture of Pluggable Databases

1.4.20 Packages

A package is an object that contains subprograms, which are programming objects that can be stored and executed in the database server, and called from other programming objects or applications. A package can contain functions or procedures, or both.

You can perform the following operations on a package by right-clicking the package name in the Connections navigator and selecting an item from the menu:

- **Edit:** Opens the package in a window, where you can modify the content and other information.
- **Run:** Lets you select a member in the package and run it.
- **Compile:** Performs a PL/SQL compilation of the members in the package.
- **Compile for Debug:** Performs a PL/SQL compilation of the members in the package, with PL/SQL library units compiled for debugging
- **Order Members By:** Orders the members of the package by location in the source, by name, or by type and by name within each type.
- **Use as Template:** Lets you create a new package using the selected package as the initial content.
- **Drop Package:** Deletes the package.
- **Create Body:** Displays a pane in which you can enter text for the package body.
- **Grant:** Lets you grant privileges on the package
- **Revoke:** Lets you revoke privileges on the package.
- **Save Package Spec and Body:** Saves the package specification and body to a file that you specify.
- **Quick DDL:** Saves the DDL statement to create the package to a file, a SQL Worksheet, or the system clipboard.

Related Topics:

[Create PL/SQL Package](#) (page 6-36)

1.4.21 Procedures

A procedure is a type of PL/SQL subprogram, which is a programming object that can be stored and executed in the database server, and called from other programming objects or applications. (Procedures do not return a value; functions return a value.)

You can perform the following operations on a procedure by right-clicking the procedure name in the Connections navigator and selecting an item from the menu:

- **Edit:** Displays the procedure text so that you can view and edit it.

- **Run:** Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the procedure in normal (not debug) mode.
- **Compile:** Performs a PL/SQL compilation of the procedure.
- **Compile for Debug:** Performs a PL/SQL compilation of the procedure, with PL/SQL library units compiled for debugging.
- **Profile** (for an Oracle Database Release 11.1 or later connection): Displays the Run/Debug/Profile PL/SQL dialog box, and then executes the procedure and collects execution statistics.
- **Drop:** Deletes the procedure.
- **Grant:** Enables you to grant available privileges on the procedure to selected users.
- **Revoke:** Enables you to revoke available privileges on the procedure from selected users.
- **Format:** Reformats the text of the procedure definition.
- **Create Unit Test:** Creates a unit test for the procedure.
- **Quick DDL:** Enables you to export the DDL statement for creating the procedure to a file, a SQL Worksheet, or the clipboard.

Related Topics:

[Run/Debug/Profile PL/SQL](#) (page 6-132)

[SQL Developer: Unit Testing](#) (page 4-1)

1.4.22 Programs

A program is an Oracle Scheduler object that describes what is to be run by the Scheduler. A program includes an action, a type, and the number of arguments that the stored procedure or external executable accepts. (A program is a separate entity from a job. A job runs at a certain time or because a certain event occurred, and invokes a certain program.)

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.23 Queues

A queue is an object in which messages are enqueued and dequeued. Queues are managed by Oracle Streams Advanced Queuing (AQ).

Related Topics:

Oracle Database Advanced Queuing User's Guide

1.4.24 Queue Tables

A queue table is a table that holds messages to be used with Oracle Streams Advanced Queuing (AQ).

Related Topics:

Oracle Database Advanced Queuing User's Guide

1.4.25 Recycle Bin

The Recycle bin (applicable only to Oracle Database Release 10g and later) holds objects that have been dropped (deleted). The objects are not actually deleted until a commit operation is performed. Before the objects are actually deleted, you can "undelete" them by selecting them in the Recycle bin and selecting **Flashback to Before Drop** from the context menu.

The Recycle bin is available only for non-system, locally managed tablespaces. Thus, to take advantage of the Recycle bin to be able to recover dropped tables for a database user, create or edit the user to have a non-system default tablespace (for example, `USERS` and not `SYSTEM`).

You can perform the following operations on an object in the Recycle bin by right-clicking the object name in the Recycle bin in the Connections navigator and selecting an item from the menu:

- **Purge:** Removes the object from the Recycle bin and deletes it.
To purge the Recycle bin, you must have the `SYSDBA` or `PURGEDBA_RECYCLEBIN` system privilege.
- **Flashback to Before Drop:** Moves the object from the Recycle bin back to its appropriate place in the Connections navigator display.

1.4.26 Replication Schemes (Oracle TimesTen In-Memory Database)

A replication scheme is a configuration, using SQL statements and a transaction-based log, whereby committed changes are copied from their source to one or more subscriber databases. The goal is to enable high efficiency and low overhead during the replication.

1.4.27 Schedules

A schedule is an Oracle Scheduler object that specifies when and how many times a job is run. Schedules can be shared by multiple jobs. For example, the end of a business quarter may be a common time frame for many jobs; instead of having to define an end-of-quarter schedule each time a new job is defined, job creators can point to a named schedule.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.28 Sequences

Sequences are used to generate unique integers. You can use sequences to automatically generate primary key values.

Related Topics:

[Create/Edit Sequence](#) (page 6-40)

Oracle Database 2 Day Developer's Guide

1.4.29 Synonyms (Public and Private)

Synonyms provide alternative names for tables, views, sequences, procedures, stored functions, packages, materialized views, Java class database objects, user-defined object types, or other synonyms. The Connections navigator has a **Synonyms** node for all synonyms (public and private) owned by the user associated with the specified connection, and a **Public Synonyms** node for all public synonyms on the database associated with the connection.

Related Topics:

[Create/Edit Synonym](#) (page 6-41)

Oracle Database 2 Day Developer's Guide

1.4.30 Tables

Tables are used to hold data. Each table typically has multiple columns that describe attributes of the database entity associated with the table, and each column has an associated data type. You can choose from many table creation options and table organizations (such as partitioned tables, index-organized tables, and external tables), to meet a variety of enterprise needs. To create a table, you can do one of the following:

- Create the table quickly by adding columns and specifying frequently used features. To do this, *do not check* the Advanced box in the Create Table dialog box.
- Create the table by adding columns and selecting from a larger set of features. To do this, *check* the Advanced box in the Create Table dialog box.
- Create the table automatically from a Microsoft Excel worksheet. To do this, right-click Tables under a connection in the Connections navigator, and select **Import Data**. When asked for the file, select a file of type .xls or .csv.

You can perform the following operations on a table by right-clicking the table name in the Connections navigator and selecting an item from the menu:

- **Edit**: Displays the Create/Edit Table (with advanced options) dialog box.
- **Import Data**: Enables you to import data from a Microsoft Excel worksheet (.xls or .csv file).
- **Export**: Enables you to export some or all of the table data to a file or to the system clipboard, in any of the following formats: XML (XML tags and data), CSV (comma-separated values including a header row for column identifiers), SQL Insert (INSERT statements), or SQL Loader (SQL*Loader control file). After you select a format, the Database Export (Unload Database Objects and Data) wizard is displayed.
- **Table**: Table actions include Rename, Copy (create a copy using a different name), Drop (delete the table), Truncate (delete existing data without affecting the table definition), Lock (set the table lock mode: row share, exclusive, and so on), Comment (descriptive comment explaining the use or purpose of the table), Parallel (change the default degree of parallelism for queries and DML on the table), No Parallel (specify serial execution), Count Rows (return the number of rows), and Generate Table API (generate statements for a PL/SQL procedure to perform operations on the table).

- **Column:** Column actions include Comment (descriptive comment about a column), Add, Drop, and Normalize.
- **Constraint:** Includes options for adding, dropping, enabling, and disabling constraints.
- **Index:** Options include Create (create an index on specified columns), Create Text (create an Oracle Text index on a column), Create Text (create a function-based index on a column), and Drop.
- **Constraint:** Options include Enable or Disable Single, Drop (delete a constraint), Add Check (add a check constraint), Add Foreign Key, and Add Unique.
- **Privileges:** If you are connected as a database user with sufficient privileges, you can Grant or Revoke privileges on the table to other users.
- **Statistics:** Options include Gather Statistics (compute exact table and column statistics and store them in the data dictionary) and Validate Structure (verifies the integrity of each data block and row, and for an index-organized table also generates the optimal prefix compression count for the primary key index on the table). Statistics are used by the Oracle Database optimizer to choose the execution plan for SQL statements that access analyzed objects.
- **Storage:** Options include Shrink Table (shrink space in a table, for segments in tablespaces with automatic segment management) and Move Table (to another tablespace). The Shrink Table options include Compact (only defragments the segment space and compacts the table rows for subsequent release, but does not readjust the high water mark and does not release the space immediately) and Cascade (performs the same operations on all dependent objects of the table, including secondary indexes on index-organized tables).
- **Trigger:** Options include Create, Create PK from Sequence (create a before-insert trigger to populate the primary key using values from a specified sequence), Enable or Disable All, Enable or Disable Single, and Drop (delete the trigger).
- **Redaction** (if the connection user has proper authorization): Enables you to add, alter, enable, disable, or drop a redaction policy on the table.
- **Enable REST Services:** Enables Oracle REST Data Services access to the selected table (available only if Oracle REST Data Services has been installed on the system containing the database associated with the connection).

You can perform the following operations on a column in a table by right-clicking the column name in the Connections navigator and selecting an item from the menu:

- **Rename:** Renames the column.
- **Drop:** Deletes the column (including all data in that column) from the table.
- **Comment:** Adds a descriptive comment about the column.
- **Encrypt** (for Oracle Database Release 10.2 and higher, and only if the Transparent Data Encryption feature is enabled for the database): Displays a dialog box in which you specify a supported encryption algorithm to be used for encrypting all data in the column. Current data and subsequently inserted data are encrypted.
- **Decrypt** (for Oracle Database Release 10.2 and higher, and only if the Transparent Data Encryption feature is enabled for the database): Decrypts data in the column

that had been encrypted, and causes data that is subsequently inserted not to be encrypted.

- **Normalize:** Creates a new table using the distinct values in the specified column. You must specify names for the new table and its primary key column, as well as a sequence name and trigger name.

The **Model** tab in a table display includes **Open in Data Modeler**, which enables you to open the table and sometimes related tables in a Data Modeler diagram view.

- If you right-click a table in the diagram and select **Show Parent and Child Tables**, any parent and child tables are added to the display (if they are not already included).
- If you right-click a table in the diagram and select **Properties**, a Data Modeler table properties display is opened in read-only mode.

Related Topics:

[Create Table \(quick creation\)](#) (page 6-41)

[Create/Edit Table \(with advanced options\)](#) (page 6-42)

[Database Export \(Unload Database Objects and Data\)](#) (page 6-77)

Oracle Database 2 Day Developer's Guide

1.4.30.1 Flashback Table Support

For tables in Oracle Database Release 11.1 and later, the table display includes the **Flashback** tab, which provides a view of the modified and original data in the table. If you have appropriate privileges, you can click the **Undo SQL** subtab to select and review the syntax required to undo changes.

Related Topics:

Oracle Database Backup and Recovery User's Guide

1.4.31 Triggers

Triggers are stored PL/SQL blocks associated with a table, a schema, or the database, or anonymous PL/SQL blocks or calls to a procedure implemented in PL/SQL or Java. Oracle Database automatically executes a trigger when specified conditions occur.

A **crossedition trigger** is intended to fire when DML changes are made in a database while an online application that uses the database is being patched or upgraded with edition-based redefinition. The body of a crossedition trigger is designed to handle these DML changes so that they can be appropriately applied after the changes to the application code are completed. To create a crossedition trigger, you must be enabled for Editions.

Related Topics:

[Create Trigger](#) (page 6-57)

[Editions](#) (page 1-16)

Oracle Database 2 Day Developer's Guide

1.4.32 Types

A data type associates a fixed set of properties with the values that can be used in a column of a table or in an argument of a function or procedure. These properties cause Oracle Database to treat values of one data type differently from values of another data type. Most data types are supplied by Oracle, although users can create data types.

Related Topics:

Oracle Database Concepts

[Create Type \(User-Defined\)](#) (page 6-59)

1.4.33 Users (Other Users)

Database users are accounts through which you can log in to the database. In the Connections navigator, you can see the **Other Users** in the database associated with a connection, but the objects that you are allowed to see for each user are determined by the privileges of the database user associated with the current database connection.

If you are connected as a user with the DBA role, you can create a database user by right-clicking Other Users and selecting **Create User**, and you can edit an existing database user by right-clicking the user under Other Users and selecting **Edit User**.

Related Topics:

[Create/Edit User](#) (page 6-59)

1.4.34 Views

Views are virtual tables (analogous to queries in some database products) that select data from one or more underlying tables. Oracle Database provides many view creation options and specialized types of views (such as materialized views) to meet a variety of enterprise needs.

Editioning views are shown in a separate navigator node if the connection is to an Oracle Database release that supports [Editions](#) (page 1-16). 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).

You can perform the following operations on a view by right-clicking the view name in the Connections navigator and selecting an item from the menu:

- **Edit:** Displays the Create/Edit View dialog box.
- **Drop:** Deletes the view.
- **Compile:** Recompile the view, to enable you to locate possible errors before runtime. You may want to recompile a view after altering one of its base tables to ensure that the change does not affect the view or other objects that depend on it.
- **Enable REST Services:** Enables Oracle REST Data Services access to the selected view (available only if Oracle REST Data Services has been installed on the system containing the database associated with the connection).

Related Topics:

[Create/Edit View](#) (page 6-62)

[Materialized Views](#) (page 1-19)

1.4.35 Window Groups

A window group is an Oracle Scheduler object that is a list of Oracle Scheduler Windows. Scheduler jobs that are scheduled to be run in a window group will be activated in that time span and using that resource plan for all windows in the group.

You can group windows for ease of use in scheduling jobs. If a job must run during multiple time periods throughout the day, week, and so on, you can create a window for each time period, and then add the windows to a window group. You can then set the `schedule_name` attribute of the job to the name of this window group, and the job executes during all the time periods specified by the windows in the window group.

For example, if you had a window called "Weekends" and a window called "Weeknights," you could add these two windows to a window group called "Downtime." The data warehousing staff could then create a job to run queries according to this Downtime window group—on weeknights and weekends—when the queries could be assigned a high percentage of available resources.

Related Topics:

[Windows](#) (page 1-27)

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.36 Windows

A window is an Oracle Scheduler object that can be used to automatically start jobs or to change resource allocation among jobs during various time periods of the day, week, and so on. A window is represented by an interval of time with a well-defined beginning and end, such as "from 12am-6am".

Windows work with job classes to control resource allocation. Each window specifies the resource plan to activate when the window opens (becomes active), and each job class specifies a resource consumer group or specifies a database service, which can map to a consumer group. A job that runs within a window, therefore, has resources allocated to it according to the consumer group of its job class and the resource plan of the window.

Related Topics:

[Scheduling Jobs Using SQL Developer](#) (page 1-67)

1.4.37 XML DB Repository

Oracle XML DB Repository is a component of Oracle Database that is optimized for handling XML data. The Oracle XML DB repository contains resources, which can be either folders (directories, containers) or files.

To create a subfolder of an existing folder, right-click the folder name and select **Create Subfolder**. To delete a folder (and optionally all subfolders under it), right-click the folder name and select **Drop Folder**.

Related Topics:

Oracle XML DB Developer's Guide

1.4.38 XML Schemas

XML schemas are schema definitions, written in XML, that describe the structure and various other semantics of conforming instance XML documents.

You can edit an XML schema by right-clicking the XML schema name in the Connections navigator and selecting **Open** from the menu; you can delete a selected schema by selecting **Drop Schema** from the menu.

Related Topics:

Oracle XML DB Developer's Guide

1.4.39 Captured and Converted Database Objects (for Migration)

If you are migrating a third-party database to Oracle, the Captured Models and Converted Models navigators can display models that include database objects, such as tables and procedures. A captured object represents an object in the captured third-party database, and a converted object represents an Oracle model of that object as it will be created in the Oracle database.

The context menu for each captured object includes **Convert to Oracle**, which creates a corresponding converted object. The context menu for each converted object includes **Generate**, which creates the corresponding Oracle Database object. (The context menus will contain other items as appropriate for the object.)

For information about the related Oracle Database objects, see the following:

- [Functions](#) (page 1-16)
- [Indexes](#) (page 1-17)
- [Procedures](#) (page 1-20)
- [Sequences](#) (page 1-22)
- [Tables](#) (page 1-23)
- [Triggers](#) (page 1-25)
- [Users \(Other Users\)](#) (page 1-26)
- [Views](#) (page 1-26)

1.5 Database Connections

A **connection** is a SQL Developer object that specifies the necessary information for connecting to a specific database as a specific user of that database. You must have at least one database connection (existing, created, or imported) to use SQL Developer.

You can connect to any target Oracle database schema using standard Oracle database authentication. Once connected, you can perform operations on objects in the database. You can also connect to schemas for MySQL and Hive, and selected third-party (non-Oracle) databases, such as Microsoft SQL Server, Sybase Adaptive Server, and IBM DB2, and view metadata and data.

When you start SQL Developer and whenever you display the database connections dialog box, SQL Developer automatically reads any connections defined in the `tnsnames.ora` file on your system, if that file exists. You can specify the `tnsnames.ora` location in the Database : Advanced preferences. By default, `tnsnames.ora` is located in the `$ORACLE_HOME/network/admin` directory, but it can also be in the directory specified by the `TNS_ADMIN` environment variable or registry value or (on Linux systems) the global configuration directory. On Windows systems, if the `tnsnames.ora` file exists but its connections are not being used by SQL Developer, define `TNS_ADMIN` as a system environment variable. For information about the `tnsnames.ora` file, see the "Local Naming Parameters (`tnsnames.ora`)" chapter in *Oracle Database Net Services Reference*.

You can create additional connections (for example, to connect to the same database but as different users, or to connect to different databases). Each database connection is listed in the Connections navigator hierarchy.

To create a new database connection, right-click the Connections node and select **New Database Connection**. Use the dialog box to specify information about the connection. You can also create a new database connection by selecting an existing connection in that dialog box, changing the connection name, changing other connection attributes as needed, and clicking Save or Connect.

To create (automatically generate) a database connection for each unlocked user account in the Oracle database instance on the local system, right-click the Connections node and select **Create Local Connections**. The connections are placed in a folder named *Auto-Generated Local Connections*. For these autogenerated connections (except for the one named `system-<database-name>`), you will always be prompted for the password when you connect, and you cannot edit the user name or password in the connection properties dialog box.

Note:

To use the Create Local Connections feature, you must specify an Oracle Home: in the Database: Advanced user preferences, click **Configure** for **Use Oracle Client**. The Oracle Home must be for Release 11.2.0.3 or later.

To edit the information about an existing database connection, right-click the connection name in the Connections navigator display and select **Properties**. Use the dialog box to modify information about the connection, but do not change the connection name.

You can use folders in the Connections navigator to organize connections into groups: for example, one folder for connections on your local system, another for connections on the test system, and another for connections on the production system.

To export information about the existing database connections into an XML file that you can later use for importing connections, right-click Connections in the Connections navigator display and select **Export Connections**. Use the dialog box to specify the connections to be exported.

To import connections that had previously been exported (adding them to any connections that may already exist in SQL Developer), right-click Connections in the Connections navigator display and select **Import Connections**. Use the dialog box to specify the connections to be imported.

To browse just the objects for a specific connection, right-click the connection name in the Connections navigator display and select **Schema Browser**. A separate Schema

Browser tab for that connection is opened in the Connections navigator, where you can click object types to view and perform operations on available objects of that type in the schema associated with the connection.

To perform limited database management operations if you are connected AS SYSDBA, right-click the connection name in the Connections navigator display and select **Manage Database**. You can click to refresh the read-only display of memory (SGA and PGA) and tablespace information. If a listener is running with a static listener configured for the database, you can also click to start and stop the database.

To perform remote debugging if you are using the Java Platform Debugger Architecture (JPDA) and you would like the debugger to listen so that a debuggee can attach to the debugger, right-click the connection name in the Connections navigator display and select **Remote Debug**. Use the dialog box to specify remote debugging information.

To estimate or compute statistics for objects in a database schema, right-click the connection name in the Connections navigator display and select **Gather Schema Statistics**. Statistics are used to optimize SQL execution.

To generate documentation in HTML format (comparable to Javadoc for Java classes) about a schema, right-click the connection name in the Connections navigator display and select **Generate DB Doc**. To view the generated documentation, open the index.html file in the output directory that you specified.

(Note that Generate DB Doc is also used for **migration projects**: right-click the migration project name in the Migration Projects navigator display and select **Migration Project Doc**; you can specify whether to include child nodes of the selected node and child master-detail per-row reports.)

To reset an expired password for the database user associated with a connection, right-click the connection name in the Connections navigator display and select **Reset Password**.

To rename a connection, right-click the connection name in the Connections navigator display and select **Rename Connection**.

To delete a connection (that is, delete it from SQL Developer, *not* merely disconnect from the current connection), right-click the connection name in the Connections navigator display and select **Delete**. Deleting a connection does not delete the user associated with that connection.

To connect using an existing connection, expand its node in the Connections navigator, or right-click its name and select **Connect**. A SQL Worksheet window is also opened for the connection. To create a separate unshared worksheet for a connection, click in the worksheet and use Ctrl+Shift+N.

To disconnect from the current connection, right-click its name in the Connections navigator and select **Disconnect**.

To reconnect to a connection where the session had been terminated, right-click its name in the Connections navigator and select **Reconnect**.

To enable Oracle REST Data Services access to objects in the schema associated with a connection, right-click its name in the Connections navigator and select **REST Services**, then **Enable REST Services** (available only if Oracle REST Data Services has been installed on the system containing the database associated with the connection). To drop support for Oracle REST Data Services access to objects in the schema associated with a connection, right-click its name in the Connections navigator and select **REST Services**, then **Drop REST Services**.

To specify a preference for using an OCI (thick, Type 2) driver (if available) instead of a JDBC (thin) driver for basic and TNS (network alias) database connections, enable the **Use OCI/Thick driver** option under the Database : Advanced user preferences.

Related Topics:

[Database](#) (page 1-99)

[Using Folders to Group Connections](#) (page 1-31)

[Export/Import Connections](#) (page 6-26)

[Debugger - Attach to JPDA](#) (page 6-79)

1.5.1 Using Folders to Group Connections

You can use folders in the Connections navigator to organize connections into groups: for example, one folder for connections on your local system, another for connections on the test system, and another for connections on the production system.

To create a folder to hold connections, right-click the name in the Connections navigator of a connection to be added to the folder, select **Add to Folder** and then **New Folder**, and specify the folder name (such as Local Connections).

To add more connections to a folder, right-click the name in the Connections navigator of a connection to be added to the folder, and select **Add to Folder** and then the name of the folder into which to add the connection.

To move a connection from one folder to another folder, right-click the connection name under its current folder, select **Add to Folder**, and then either the name of the destination folder or **New Folder** to move the connection to a new folder to be created.

To remove a connection from the folder, right-click the connection name under the folder and select **Remove from Folder**. (This does *not* delete the connection; it is moved to the top level in the Connections navigator hierarchy display.)

To remove a folder, right-click the folder name select **Remove Folder**. (This does *not* delete any connections that are in the folder; these connections are moved to the top level in the Connections navigator hierarchy display.)

To rename a folder, right-click the folder name, select **Rename Folder**, and specify the new name.

1.5.2 Sharing of Connections

By default, each connection in SQL Developer is shared when possible. For example, if you open a table in the Connections navigator and two SQL Worksheets using the same connection, all three panes use one shared connection to the database. In this example, a commit operation in one SQL Worksheet commits across all three panes. If you want a dedicated session, you must duplicate your connection and give it another name. Sessions are shared by name, not connection information, so this new connection will be kept separate from the original.

1.5.3 Advanced Security for JDBC Connection to the Database

You are encouraged to use Oracle Advanced Security to secure a JDBC or OCI connection to the database. Both the JDBC OCI and the JDBC Thin drivers support at least some of the Oracle Advanced Security features. If you are using the OCI driver, you can set relevant parameters in the same way that you would in any Oracle client setting. The JDBC Thin driver supports the Oracle Advanced Security features

through a set of Java classes included with the JDBC classes in a Java Archive (JAR) file and supports security parameter settings through Java properties objects.

Related Topics:

Oracle Database JDBC Developer's Guide

1.5.4 Connections with Operating System (OS) Authentication

When you create a connection to an Oracle database that is using operating system (OS) authentication, you can omit the user name and password; that is, specify a connection name and all the other necessary information, except do not specify a user name or password.

If you omit the user name and password trying to create a connection to a system that is not configured for external authentication, an error message is displayed.

See Also:

For information about using external authentication, including the use of the OS_AUTHENT_PREFIX and REMOTE_OS_AUTHENT database initialization parameters, see *Oracle Database Security Guide*

1.5.5 Connections with Proxy Authentication

Proxy authentication enables one JDBC connection to act as a proxy for other JDBC connections. If you use the Proxy Connection option when you create a database connection, the connection will be used to connect as the specified user for the connection, but authenticated using the user name and either the password or distinguished name of the proxy user. For information about using a middle tier server for proxy authentication, see *Oracle Database Security Guide*.

For example, to create connection for a user named PROXY_USER but connecting using the user name and password of existing database user SCOTT, follow these steps.

1. Create the proxy user and grant it the appropriate privileges:

```
CREATE USER proxy_user IDENTIFIED BY <password>;
ALTER USER proxy_user GRANT CONNECT THROUGH scott AUTHENTICATED USING PASSWORD;
GRANT create session TO proxy_user;
. . .<Grant other privileges as needed.>
```

2. Create a new database connection. For example: connection name = proxy_conn, user name = scott, password = *<password for scott>*.
3. Enable (check) **Proxy Connection**.
4. In the Oracle Proxy Connection dialog box, select **User Name** for Proxy Type
5. For Proxy User, enter **PROXY_USER**; and for Proxy Password, enter the password for the PROXY_USER database user.
6. Click OK to close the Oracle Proxy Connection dialog box.
7. Complete any other necessary connection information, and click Connect to create the connection.

In this example, when you connect using the proxy_conn connection, the user name and password for user SCOTT are used to connect to the database, but the connection sees those objects that the PROXY_USER user is permitted to see.

The preceding instructions cause two sessions to be started (one for the proxy user and one for the proxy client) when the connection is opened. If you want to have a single session (with no second password or distinguished name required) started, you can follow these steps instead after you create the proxy user and grant it appropriate privileges:

1. Create a new database connection. For example: connection name = proxy_conn, user name = proxy_user[scott], password = <password for proxy_user>.
2. For Connection Type, specify **TNS**.
(Note: Do *not* enable Proxy Connection.)
3. For Network Alias, select the network alias for the database for the connection.
4. Complete any other necessary connection information, and click **Connect** to create the connection.

1.5.6 Connections with SSH Authentication

SSH (Secure Shell) authentication enables you to specify that the SSH protocol should be used to authenticate users attempting to connect to the database. To enable SSH authentication, click **View**, and then click **SSH**. In the SSH Hosts pane, right-click and select **New SSH Host**.

Name: A name for the SSH session.

Host: SSH server. SQL Developer will create an SSH session to this host, using the specified details.

Port: SSH port. The default port is 22.

Username: User name that will be used to authorize the SSH session.

Use Key File: Specifies that a key file should be used to provide authentication. The key file contains a private key that should correspond to a public key registered with the server. The key file must be in OpenSSH format. The server verifies that SQL Developer has access to the proper private key and thus the user is who he or she claims to be.

Key File: Path to the key file.

When you create an SSH (Secure Shell), you can create one or more local port forwards and/or remote port forwards.

Name: A name to be given to the local or remote port forward.

Host: Host name for the forward.

Port: Port on the specified host.

Automatically assign [local | remote] port: Lets a port be assigned automatically.

Use specific [local | remote] port: Specify the port to be used.

For detailed information about on the SSH protocol and authorization, see the OpenBSD man page for ssh: <http://www.openbsd.org/cgi-bin/man.cgi?query=ssh&sektion=1>

You may also find the following article helpful: "SSH Tunnel - Local and Remote Port Forwarding Explained With Examples" (<http://blog.sensible.io/2014/05/17/ssh-tunnel-local-and-remote-port-forwarding-explained-with-examples.html>)

When using SSH authentication, note the following:

- The host name entered into the main connection properties is resolved by the SSH host. So, to connect to a database on the same host, the SSH host is the remote server name and the connection host is **localhost**.
- If you choose not to use a key file, or if the selected key file has an associated passphrase (an added layer of security), a user attempting to open the connection will be prompted for the password or passphrase.

1.6 Entering and Modifying Data

You can use SQL Developer to enter data into [Tables](#) (page 1-23) and to edit and delete existing table data. To do any of these operations, select the table in the Connections navigator, then click the **Data** tab in the table detail display. (If you click one of the other tabs -- **Columns**, **Constraints**, **Grants**, **Statistics**, **Triggers**, and so on -- you can view the information and perform relevant context menu (right-click) operations.)

The following figure shows the Data pane for a table named BOOKS, with a filter applied to show only books whose rating is 10, and after the user has double-clicked in the Title cell for the first book.

BOOK_ID	TITLE	AUTHOR_LAST_NAME	AUTHOR_FIRST_NAME	RATING
1	A1111 Moby Dick	Melville	Herman	10
2	A5555 Software Wizardry	Abugov	D.	10

Icons and other controls under the Data tab provide the following options:

- **Freeze View** (the **pin**) keeps that object's tab and information in the window when you click another object in the Connections navigator; a separate tab and display are created for that other object. If you click the pin again, the object's display is available for reuse.
- **Refresh** queries the database to update the data display. If a filter is specified, the refresh operation uses the filter.
- **Insert Row** adds an empty row after the selected row, for you to enter new data.
- **Delete Selected Row(s)** marks the selected rows for deletion. The actual deletion does not occur until you commit changes.
- **Commit Changes** ends the current transaction and makes permanent all changes performed in the transaction.
- **Rollback Changes** undoes any work done in the current transaction.

- **Sort** displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.
- **Filter** enables you to enter a SQL predicate (WHERE clause text without the WHERE keyword) for limiting the display of data. For example, to show only rows where the RATING column value is equal to 10, specify rating=10 and press Enter. (To see any previously applied filters, clear the Filter box, then either press Ctrl + spacebar or click the drop-down arrow in the Filter box.)
- **Filter Column** enables you to enter a partial value (such as a number or a string; at least two characters for a string), to limit the dialog box display to items containing the partial value, so that you can then select the one item to appear in the grid. For example, entering EMP for column names might show a list of columns including EMPLOYEE_ID and IS_TEMP.
- **Actions** displays a menu with actions relevant to the table.

When you enter a cell in the grid, you can directly edit the data for many data types, and for all data types you can click the ellipsis (...) button to edit the data. For binary data you cannot edit the data in the cell, but must use the ellipsis button.

In the data grid, the context menu (right-click) includes the following commands:

- **Single Record View** displays the Single Record View dialog box, which enables you to edit data for a table or view, one record at a time.
- **Auto-fit All Columns** adjusts the width of all columns according to your specification (by column header, by column data, or best fit).
- **Auto-fit Selected Columns** adjusts the width of the selected columns according to your specification (by column header, by column data, or best fit).
- **Count Rows** displays the number of rows in the table.
- **Find/Highlight** displays the Find/Highlight, which enables you to find, and optionally highlight, data value matches in the table data grid.
- **Publish to Apex** (if Application Express is installed) creates a small Application Express application based on the data. It displays a dialog box in which you specify the following for the application to be created: workspace, application name, theme, page name, and SQL statement for generating the report.
- **Publish to REST** displays the RESTful Services wizard, which enables you to publish the query associated with the grid to REST to create the resource module, template, and handler.
- **Export** enables you to export some or all of the table data to a file or to the system clipboard, in any of the following formats: XML (XML tags and data), CSV (comma-separated values including a header row for column identifiers), SQL Insert (INSERT statements), or SQL Loader (SQL*Loader control file). After you select a format, the Database Export (Unload Database Objects and Data) wizard is displayed.

You can copy and paste data between table grid cells and cells in a Microsoft Excel worksheet.

To copy table data to the clipboard, click the column header (for all column data) or select specific cells and press Ctrl+C; to copy the column header text along with the table data, press Ctrl+Shift+C.

To sort the display of rows by values within a column, double-click the column header; to switch between ascending and descending sorting, double-click the up/down arrow in the column header.

In the Data pane for a table or view, you can **split** the display vertically or horizontally to see two (or more) parts independently by using the split box (thin blue rectangle), located to the right of the bottom scroll bar and above the right scroll bar.

In the Data pane, the acceptable format or formats for entering dates may be different from the date format required by SQL*Plus.

Related Topics:

[Single Record View](#) (page 6-138)

[Find/Highlight](#) (page 6-92)

[Database Export \(Unload Database Objects and Data\)](#) (page 6-77)

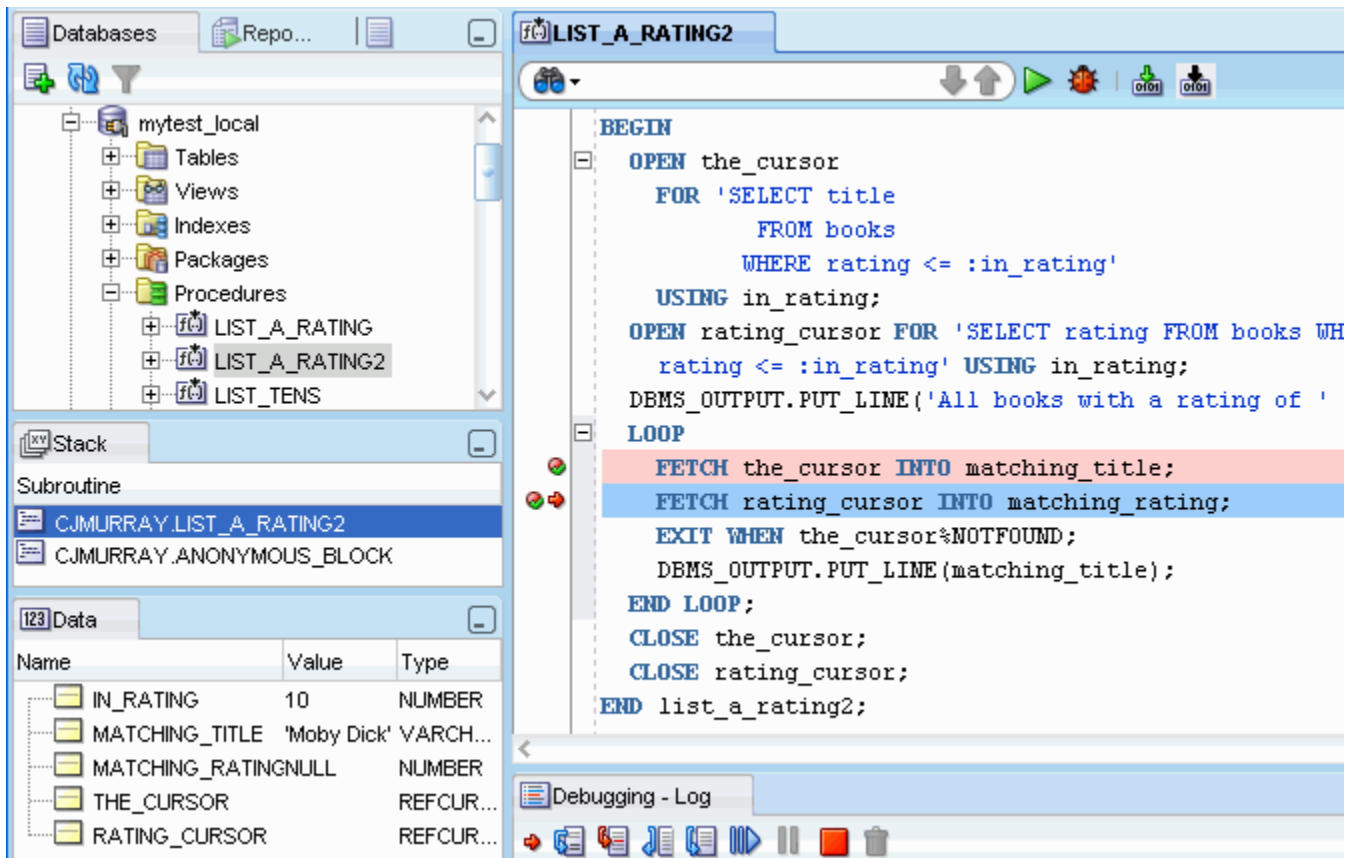
[Publish to REST](#) (page 6-122)

1.7 Running and Debugging Functions and Procedures

You can use SQL Developer to run and debug PL/SQL subprograms (functions and procedures).

- To run a subprogram, click its name in the Connections navigator; then either right-click and select Run, or click the Edit icon and then click the Run icon above its source listing.
- To debug a subprogram, click its name in the Connections navigator. If the procedure in its current form has not already been compiled for debug, right-click and select Compile for Debug. Then click the Edit icon and click the Debug icon above its source listing.

In both cases, a code editing window is displayed. The following figure shows the code editing window being used to debug a procedure named LIST_A_RATING2, which is used for tutorial purposes in [Debug a PL/SQL Procedure](#) (page 5-10).



In the code editing window, under the tab with the name of the subprogram, is a toolbar, and beneath it is the text of the subprogram, which you can edit. You can set and unset breakpoints for debugging by clicking to the left of the thin vertical line beside each statement with which you want to associate a breakpoint. (When a breakpoint is set, a red circle is displayed.)


The toolbar under the tab for the subprogram name includes the icons shown in the following figure.



- **Freeze Content** (the **pin**) keeps that subprogram's tab and information in the window when you click another object in the Connections navigator; a separate tab and display are created for that other object. If you click the pin again, the object's display is available for reuse.
- **Run** starts normal execution of the subprogram, and displays the results in the **Running - Log** tab.
- **Debug** starts execution of the subprogram in debug mode, and displays the **Debugging - Log** tab, which includes the debugging toolbar for controlling the execution.
- **Compile for Debug** performs a PL/SQL compilation of the subprogram so that it can be debugged.
- **Compile** performs a PL/SQL compilation of the subprogram.

- **Switch to Read Only / Switch to Write Mode** toggles between read-only and read/write mode for the code editing window. (See also the [Code Editor](#) (page 1-94) user preference Start in Read Only Mode.)
- **Profile** displays the [Run/Debug/Profile PL/SQL](#) (page 6-132) dialog box.

The **Debugging - Log** tab under the code text area contains the debugging toolbar and informational messages. The debugging toolbar has the icons shown in the following figure.



- **Find Execution Point** goes to the execution point (the next line of source code to be executed by the debugger).
- **Step Over** bypasses the next subprogram (unless the subprogram has a breakpoint) and goes to the next statement after the subprogram. If the execution point is located on a subprogram call, it runs that subprogram without stopping (instead of stepping into it), then positions the execution point on the statement that follows the call. If the execution point is located on the last statement of a subprogram, Step Over returns from the subprogram, placing the execution point on the line of code that follows the call to the subprogram from which you are returning.
- **Step Into** executes a single program statement at a time. If the execution point is located on a call to a subprogram, Step Into steps into that subprogram and places the execution point on its first statement. If the execution point is located on the last statement of a subprogram, Step Into returns from the subprogram, placing the execution point on the line of code that follows the call to the subprogram from which you are returning.
- **Step Out** leaves the current subprogram and goes to the next statement.
- **Step to End of Method** goes to the last statement of the current subprogram.
- **Resume** continues execution.
- **Pause** halts execution but does not exit, thus allowing you to resume execution.
- **Terminate** halts and exits the execution. You cannot resume execution from this point; instead, to start running or debugging from the beginning of the subprogram, click the Run or Debug icon in the Source tab toolbar.

The **Breakpoints** tab displays breakpoints, both system-defined and user-defined.

The **Smart Data** tab displays information about variables, using your Debugger: Smart Data preferences. You can also specify these preferences by right-clicking in the Smart Data window and selecting **Preferences**.

The **Data** tab displays information about variables, using your Debugger: Data preferences. You can also specify these preferences by right-clicking in the Data window and selecting **Preferences**.

The **Watches** tab displays information about watches.

If the function or procedure to be debugged is on a remote system, see also [Remote Debugging](#) (page 1-39).

1.7.1 Using Bookmarks When Editing Functions and Procedures

When you are editing a long function or procedure, you may find it convenient to create bookmarks in the code so that you can easily navigate to points of interest.

To create or remove a bookmark, click **Navigate**, then **Toggle Bookmark**. When a bookmark is created, an icon appears to the left of the thin vertical line.

To go to a specific bookmark, click **Navigate**, then **Go to Bookmark**. To go to the next or previous bookmark, click **Navigate**, then **Go to Next Bookmark** or **Go to Previous Bookmark**, respectively.

To remove all bookmarks from the currently active editing window for a function or procedure or from all open editing windows, click **Navigate**, then **Remove Bookmarks from File** or **Remove All Bookmarks**, respectively.

You can also go to a specific line or to your last edit by clicking **Navigate**, then **Go to Line** or **Go to Last Edit**, respectively.

1.7.2 Remote Debugging

Remote debugging allows you to debug a procedure that is initiated from a program other than SQL Developer. You can also use remote debugging to debug stored procedures initiated from computers other than your local system.

Remote debugging involves many of the steps as for local debugging; however, you need to set up the SQL Developer debugger to listen. Do the following before you start the remote debugging:

1. Do one of the following:

- Use an Oracle client such as SQL*Plus to issue the debugger connection command. Whatever client you use, make sure that the session which issues the debugger connection commands is the same session which executes your PL/SQL program containing the breakpoints (and be sure each breakpoint is an executable line, not a declaration). For example, if the name of the remote system is `remote1`, use the following SQL*Plus statement to open a TCP/IP connection to that system and the port for the JDWP session:

```
EXEC DBMS_DEBUG_JDWP.CONNECT_TCP('remote1', '4000');
```

The first parameter is the IP address or host name of the remote system, and the second parameter is the port number on that remote system on which the debugger is listening.

- Define the `ORA_DEBUG_JDWP` operating system environment variable. For example:

```
ORA_DEBUG_JDWP=host=mypc:port=1234
```

2. Right-click the connection for the remote database, select **Remote Debug**, and complete the information in the Debugger — Attach to JPDA dialog box.

Then, follow the steps that you would for local debugging.

Related Topics:

[Debugger - Attach to JPDA](#) (page 6-79)

[Debug a PL/SQL Procedure](#) (page 5-10)

1.7.3 Displaying SQL Trace (.trc) Files

If you have any SQL Trace (.trc) output files, you can display them in SQL Developer as an alternative to using the TKPROF program to format the contents of the trace file. To open a .trc file in SQL Developer and see an attractive, effective display of the information, click **File**, then **Open**, and specify the file; or drag the file's name or icon into the SQL Developer window.

You can then examine the information in the List View, Statistics View, and History panes, with each pane including options for filtering and controlling the display.

Related Topics:

Oracle Database SQL Tuning Guide

1.7.4 Using the PL/SQL Hierarchical Profiler

For an Oracle Database Release 11.1 or later connection, you can use the PL/SQL hierarchical profiler to identify bottlenecks and performance-tuning opportunities in PL/SQL applications. Profiling consists of the two steps: running the PL/SQL module in profiling mode, and analyzing the reports. In addition, some one-time setup work is required the first time you use profiling in SQL Developer.

To initiate profiling, right-click the name of the function or procedure in the Connections navigator hierarchy and select **Profile**, or click the Profile button on the PL/SQL source editor toolbar. After the function or procedure is run in profiling mode, the profiler reports are located at the Execution Profiles tab of the object viewer window. You can review 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, including all subprograms that a given subprogram called (that is, children of the given subprogram)

Related Topics:

Oracle Database Development Guide

1.7.5 Setting Expression Watches

A watch enables you to monitor the changing values of variables or expressions as your program runs. After you enter a watch expression, the Watches window displays the current value of the expression. As your program runs, the value of the watch changes as your program updates the values of the variables in the watch expression.

A watch evaluates an expression according to the current context which is controlled by the selection in the Stack window. If you move to a new context, the expression is reevaluated for the new context. If the execution point moves to a location where any of the variables in the watch expression are undefined, the entire watch expression becomes undefined. If the execution point returns to a location where the watch expression can be evaluated, the Watches window again displays the value of the watch expression.

To open the Watches window, click **View**, then **Debugger**, then **Watches**.

To add a watch, right-click in the Watches window and select **Add Watch**. To edit a watch, right-click in the Watches window and select **Edit Watch**.

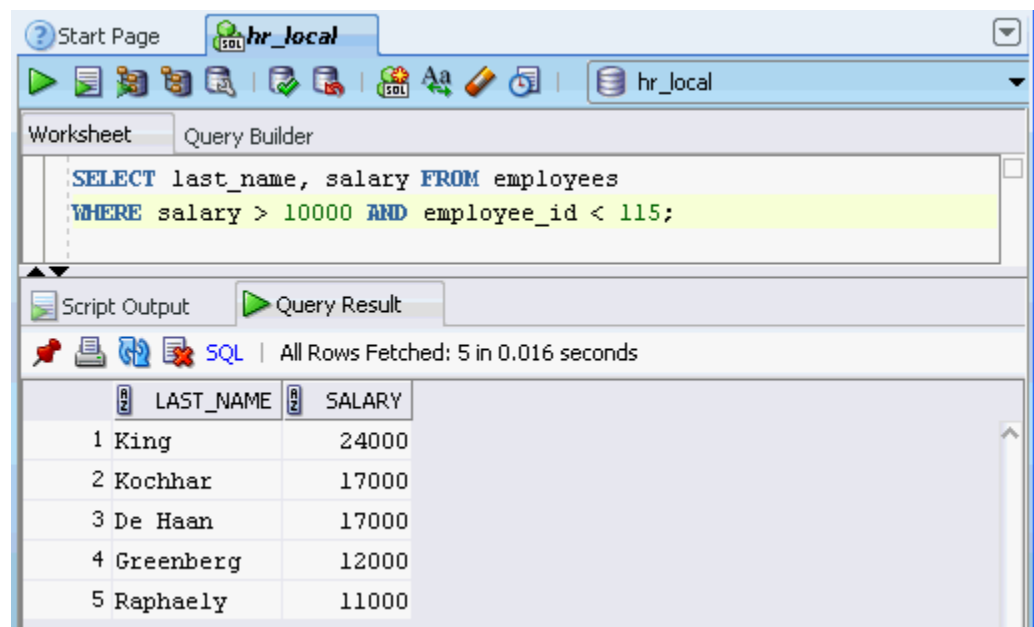
1.8 Using the SQL Worksheet

You can use the SQL Worksheet to enter and execute SQL, PL/SQL, and SQL*Plus statements. You can specify any actions that can be processed by the database connection associated with the worksheet, such as creating a table, inserting data, creating and editing a trigger, selecting data from a table, and saving that data to a file.

You can display a SQL Worksheet by right-clicking a connection in the Connections navigator and selecting **Open SQL Worksheet**, by selecting **Tools** and then **SQL Worksheet**, or by clicking the **Use SQL Worksheet** icon under the menu bar. In the Select Connection dialog box, select the database connection to use for your work with the worksheet. You can also use that dialog box to create and edit database connections. (You can have a SQL Worksheet window open automatically when you open a database connection by enabling the appropriate SQL Developer user preference under Database Connections.)

To create a separate unshared worksheet for a connection, click in the worksheet and use Ctrl+Shift+N.

The SQL Worksheet has Worksheet and Query Builder tabs, as shown in the following figure (where the Worksheet tab is selected):



SQL Worksheet toolbar (under the Worksheet tab): Contains icons for the following operations:

- **Execute Statement** executes the statement at the mouse pointer in the Enter SQL Statement box. The SQL statements can include bind variables and substitution variables of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary); a pop-up box is displayed for entering variable values.

- **Run Script** executes all statements in the Enter SQL Statement box using the Script Runner. The SQL statements can include substitution variables (but not bind variables) of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary); a pop-up box is displayed for entering substitution variable values.

Every time you click the Run Script icon, the linesize value is reset to the system default, which is the width of the Script Output pane. If you want to ensure a specific linesize is in effect, use the `set linesize` command (for example, `set linesize=80`) for each Run Script occurrence.
- **Commit** writes any changes to the database, and ends the transaction; also clears any output in the Results and Script Output panes.
- **Rollback** discards any changes without writing them to the database, and ends the transaction; also clears any output in the Results and Script Output panes.
- **Cancel** stops the execution of any statements currently being executed.
- **Monitor SQL Status** (Oracle Database Release 11.1 and later only) calls the real-time SQL monitoring feature of Oracle Database, enabling you to monitor the performance of SQL statements while they are executing.
- **Explain Plan** generates the execution plan for the statement (internally executing the EXPLAIN PLAN statement). To see the execution plan, click the Explain tab.
- **SQL Tuning Advisor** displays a window that runs the SQL Tuning Advisor, which is SQL diagnostic software in the Oracle Database Tuning Pack. The Overview and Details tabs include advice or recommendations for how to tune the specified statement, along with a rationale and expected benefit.
- **Autotrace** generates trace information for the statement. To see the Autotrace Pane, click the Autotrace tab.
- **Clear** erases the statement or statements in the Enter SQL Statement box.
- **Save REST Handler** saves the REST Handler definitions directly to the database. This icon is available only when adding or editing a resource handler from the Connections navigator.
- To the right of these icons is a drop-down list for changing the **database connection** to use with the worksheet.

The context menu (right-click, or Shift+F10) includes the preceding Worksheet toolbar operations, plus the following operations (some depending on the type of object displayed in the worksheet):

- **Print File** prints the contents of the Enter SQL Statement box.
- **Cut, Copy, Paste, and Select All** have the same meanings as for normal text editing operations.
- **Refactoring** enables you to do the following on selected text: switch character case (to upper/lower/initcap), extract the sequence of PL/SQL statements to a procedure, or rename the local variable.
- **Format** formats the SQL statement (capitalizing the names of statements, clauses, keywords, and so on). The Database: SQL Formatter preferences are used for the formatting. (You can also use the command-line interface for formatting.)

- **Advanced Format** displays a dialog box where you can specify the output destination and output type for the formal operation.
- **Quick Outline** displays the Outline pane with a graphical outline of the object displayed in the worksheet (if an outline is relevant for this type of object). You can click a node in the outline to go to the associated place in the text in the worksheet.
- **Popup Describe**, if the name of a database object is completely selected, displays a window with tabs and information appropriate for that type of object.
- **Save Snippet** opens the Save Snippet (User-Defined) dialog box with the selected text as the snippet text.

Enter SQL Statement: The statement or statements that you intend to execute. For multiple statements, each non-PL/SQL statement must be terminated with either a semicolon or (on a new line) a slash (/), and each PL/SQL statement must be terminated with a slash (/) on a new line. SQL keywords are automatically highlighted. To format the statement, right-click in the statement area and select **Format SQL**.

You can drag some kinds of objects from the Connections navigator and drop them into the Enter SQL Statement box:

- If you drag and drop a table or view, by default a SELECT statement is constructed with all columns in the table or view. You can then edit the statement, for example, modifying the column list or adding a WHERE clause.
- If you drag and drop a function or procedure, a snippet-like text block is constructed for you to edit when including that object in a statement.

To view details for any object, you can select its name in the Enter SQL Statement box and select **Popup Describe** from the context menu (or press Shift+F4). For example, if you select a table name and press Shift+F4, information about Columns, Constraints, Grants, and so on is displayed; or if you select a procedure name and press Shift+F4, information about Code, Grants, Dependencies, References, and Details is displayed.

Tabs display panes with the following information:

- **Results:** Displays the results of the most recent Execute Statement operation.
- **Explain:** Displays the output if you clicked the Explain Execution Plan icon.
- **Script Output:** Displays the output if you clicked the Run Script icon.
- **DBMS Output:** Displays the output of DBMS_OUTPUT package statements.
- **OWA Output:** Displays Oracle Web Agent (MOD_PLSQL) output.

To toggle the relative heights of the Enter SQL Statement area and the area for tabs and display panes, press Ctrl+Alt+L. You can also manually adjust the heights.

Related Topics:

[Execution Plan](#) (page 1-46)

[OWA Output Pane](#) (page 1-47)

[DBMS Output Pane](#) (page 1-46)

1.8.1 SQL*Plus Statements Supported and Not Supported in SQL Worksheet

The SQL Worksheet supports SQL*Plus statements. SQL*Plus statements must be interpreted by the SQL Worksheet before being passed to the database; any statements that are not supported by the SQL Worksheet are ignored and not passed to the database.

The following SQL*Plus statements are supported by the SQL Worksheet:

```
@
@@
/
acc[ept]
autotrace
clear screen
conn[ect]
def[ine]
desc[ribe]
doc[ument]
echo
errors
esc[ape]
exec[ute]
exit
feed[back]
help
ho[st]
pagesize
pau[se]
print
pro[mpt]
quit
rem[ark]
roll[back]
show (SHOW ALL and individual options shown in SHOW ALL output)
set (valid SET options corresponding to options in SHOW ALL output)
spo[ol] {filename[.ext] | OFF}
sta[rt]
term[out]
timi[ng]
undef[ine]
user
var[iable]
ver[ify]
whenever
xquery
```

The following statements are *not* supported by the SQL Worksheet:

```
a[ppend]
archive
attr[ibute]
bre[ak]
bti[tile]
c[hange]
comp[ute]
copy
del
disc[onnect]
ed[it]
```



```

get
i[nput]
l[ist]
newpage
oradebug
passw[ord]
r[un]
recover
repf[ooter]
reph[earer]
sav[e]
set/show markup
set/show sqlnumber
startup
shu[tdown]
startup
store
tti[tile]

```

For information about SQL*Plus statements, you can enter the `help` statement. For information about a specific statement or topic, include it as the parameter (for example, `help @`, `help exit`, or `help reserved words`). If the statement is not supported in SQL Developer, or if there are restrictions or usage notes, the help display includes this information (for example, "Not currently in SQL Developer").

1.8.2 Script Runner

The script runner emulates a limited set of SQL*Plus features. You can often enter SQL and SQL*Plus statements and execute them by clicking the **Run Script** icon. The Script Output pane displays the output.

The SQL*Plus features available in the script runner include `@`, `@@`, `CONNECT`, `EXIT`, `QUIT`, `UNDEFINE`, `WHENEVER`, and substitution variables. For example, to run a script named `c:\myscripts\mytest.sql`, type `@c:\myscripts\mytest` in the Enter SQL Statement box, and click the drop-down next to the Execute Statement icon and select Run Script.

The following considerations apply to using the SQL Developer script runner:

- You cannot use bind variables. (The Execute SQL Statement feature does let you use bind variables of type `VARCHAR2`, `NUMBER`, and `DATE`.)
- For substitution variables, the syntax `&&variable` assigns a permanent variable value, and the syntax `&variable` assigns a temporary (not stored) variable value.
- For `EXIT` and `QUIT`, `commit` is the default behavior, but you can specify `rollback`. In either case, the context is reset: for example, `WHENEVER` command information and substitution variable values are cleared.
- `DESCRIBE` works for most, but not all, object types for which it is supported in SQL*Plus.
- For SQL*Plus statements that are not supported, a warning message is displayed.
- SQL*Plus comments are ignored.
- For XMLType data, data in the column is displayed as "SYS.XMLDATA" if the database connection uses a JDBC Thin driver, but the expanded XML values are displayed if the connection uses an OCI (thick, Type 2) driver.

If you have SQL*Plus available on your system, you may want to use it instead of the script runner.

1.8.3 Execution Plan

The Execute Explain Plan icon generates the execution plan, which you can see by clicking the Explain Plan tab. The execution plan is the sequence of operations that will be performed to execute the statement. An execution plan shows a row source tree with the hierarchy of operations that make up the statement. For each operation, it shows the ordering of the tables referenced by the statement, access method for each table mentioned in the statement, join method for tables affected by join operations in the statement, and data operations such as filter, sort, or aggregation.

In addition to the row source tree, the plan table displays information about optimization (such as the cost and cardinality of each operation), partitioning (such as the set of accessed partitions), and parallel execution (such as the distribution method of join inputs).

1.8.4 Autotrace Pane

The Autotrace pane displays trace-related information when you execute the SQL statement by clicking the **Autotrace** icon. Most of the specific information displayed is determined by the SQL Developer Preferences for Database: Autotrace/Explain Plan. If you cancel a long-running statement, partial execution statistics are displayed.

This information can help you to identify SQL statements that will benefit from tuning. For example, you may be able to optimize predicate handling by transitively adding predicates, rewriting predicates using Boolean algebra principles, moving predicates around in the execution plan, and so on.

To use the autotrace feature, the database user for the connection must have the `SELECT_CATALOG_ROLE` and `SELECT ANY DICTIONARY` privileges.

Related Topics:

[SQL Developer Preferences](#) (page 1-92)

1.8.5 DBMS Output Pane

The PL/SQL `DBMS_OUTPUT` package enables you to send messages from stored procedures, packages, and triggers. The `PUT` and `PUT_LINE` procedures in this package enable you to place information in a buffer that can be read by another trigger, procedure, or package. In a separate PL/SQL procedure or anonymous block, you can display the buffered information by calling the `GET_LINE` procedure. The DBMS Output pane is used to display the output of that buffer. This pane contains icons and other controls for the following operations:

- **Add New DBMS Output Tab:** Prompts you to specify a database connection, after which a tab is opened within the DBMS Output pane for that connection, and the `SET SERVEROUTPUT` setting is turned on so that any output is displayed in that tab. (To stop displaying output for that connection, close the tab.)
- **Clear:** Erases the contents of the pane.
- **Save:** Saves the contents of the pane to a file that you specify.
- **Print:** Prints the contents of the pane.

- **Buffer Size:** For databases before Oracle Database 10.2, limits the amount of data that can be stored in the DBMS_OUTPUT buffer. The buffer size can be between 1 and 1000000 (1 million).

1.8.6 OWA Output Pane

OWA (Oracle Web Agent) or MOD_PLSQL is an Apache (Web Server) extension module that enables you to create dynamic web pages from PL/SQL packages and stored procedures. The OWA Output pane enables you to see the HTML output of MOD_PLSQL actions that have been executed in the SQL Worksheet. This pane contains icons for the following operations:

- **Add New OWA Output Tab:** Prompts you to specify a database connection, after which a tab is opened within the OWA Output pane for that connection, and entries written to the OWA output buffer are displayed in that tab. (To stop displaying output for that connection, close the tab.)
- **Clear:** Erases the contents of the pane.
- **Save:** Saves the contents of the pane to a file that you specify.
- **Print:** Prints the contents of the pane.

1.8.7 SQL History

You can click **View**, then **History** (or press **F8** in the SQL Worksheet) to view a dockable window with SQL statements and scripts that you have executed, and optionally select one or more statements to have them either replace the statements currently on the SQL Worksheet or be added to the statements currently on the SQL Worksheet.

You can click on a column heading to sort the rows by the values in that column.

The SQL history list will not contain any statement that can include a password. Such statements include (but are not necessarily limited to) CONNECT, ALTER USER, and CREATE DATABASE LINK.

You can control the maximum number of statements in the history by setting the SQL History Limit preference (see [Database](#) (page 1-99): Worksheet preferences).

Append: Appends the selected statement or statements to any statements currently on the SQL Worksheet. You can also append the selected statement or statements by dragging them from the SQL History window and dropping them at the desired location on the SQL Worksheet.

Replace: Replaces any statements currently on the SQL Worksheet with the selected statement or statements.

Clear History: Removes the selected statement or statements (or all statements if no statements are selected) from the SQL history. (You will be asked to confirm this action.)

Filter: If you type a string in the text box and click Filter, only SQL statements containing that string are displayed.

1.8.8 Query Builder

The Query Builder tab in the SQL Worksheet enables you to display and build SQL queries graphically. You can create a SELECT statement by dragging and dropping table and view names and by graphically specifying columns and other elements of

the query. While you are building the query, you can click the Worksheet tab to see the SELECT statement reflecting current specifications, and then click the Query Builder tab to continue building the query if you want.

In the area below the graphical display of tables and views, you can specify one or more lines with the following information:

Output: Specifies whether to include the expression in the statement output.

Expression: Column name or expression.

Aggregate: Aggregation function to be used (Avg, Avg Distinct, Count, and so on).

Alias: Column alias to be used.

Sort Type: Ascending or Descending sorting of results.

Sort Order: Order to use in sorting results if multiple columns or expressions are to be used (for example, sorting first by department and then by salary within each department).

Grouping: Specifies whether to insert a GROUP BY clause.

Criteria: An expression with one or more criteria that must be satisfied for a result to be returned. You can specify any WHERE clause (without the WHERE keyword). For example, for `employees.SALARY`, specifying `> 10000` limits the results to employees with salaries greater than \$10,000.

Or: You can specify one or more OR clauses to be added to the query criteria. For example, if `Expression = employees.LAST_NAME`, you could specify `Or as = 'Smith'` to add `OR (employees.LAST_NAME = 'Smith')` to the query.

1.8.9 Command-Line Interface for SQL Formatting

As an alternative to using the SQL Developer graphical interface for formatting a .sql file or all .sql files in a directory or folder, you can use the command line.

Before invoking the command-line interface for SQL formatting, start the SQL Developer graphical interface, so that the Database: SQL Formatter preferences (which are used for the formatting) are loaded and available.

The following example takes the SQL code in `c:\temp\myfile.sql` and creates `c:\temp\myfile_out.sql` containing the formatted code. (Enter the command on one command line.)

```
C:\sqldeveloper\sqldeveloper\bin>sdcli format input=c:\temp\myfile.sql output=c:\temp\myfile_out.sql
```

Related Topics:

[Command-Line Interface for SQL Developer](#) (page 1-123)

1.8.10 SQL Worksheet "Hints" for Formatting Output

You can use special SQL Worksheet "hints" to generate output in several formats, such as CSV and SQL INSERT statements. (These hints do not work in SQL*Plus but do work in SQLcl.) You must use Run Script (F5), not Execute Statement, to see the formatted output. The hints must be in lowercase. Some example statements showing the available special SQL Worksheet hints are:

```
SELECT /*ansiconsole*/ * FROM books; -- Best appearance for ANSI terminal display
SELECT /*csv*/ * FROM books; -- Comma-separated values
SELECT /*delimited*/ * FROM books; -- (same as csv)
```

```

SELECT /*fixed*/ * FROM books; -- Fixed-width fields with trailing blanks
SELECT /*html*/ * FROM books; -- Marked-up HTML table
SELECT /*insert*/ * FROM books; -- SQL INSERT statements
SELECT /*json*/ * FROM books; -- JSON object format
SELECT /*loader*/ * FROM books; -- Pipe-delimited format suitable for SQL*Loader
SELECT /*text*/ * FROM books; -- Plain text
SELECT /*xml*/ * FROM books; -- Tagged XML

```

The following example shows the output generated by the first statement (`SELECT /*csv*/ * FROM books;`). It assumes that the BOOKS table was created and populated using the instructions in [SQL Developer Tutorial: Creating Objects for a Small Database](#) (page 5-1).

```

"BOOK_ID","TITLE","AUTHOR_LAST_NAME","AUTHOR_FIRST_NAME","RATING"
"A1111","Moby Dick","Melville","Herman",10
"A2222","Get Rich Really Fast","Scammer","Ima",1
"A3333","Finding Inner Peace","Blissford","Serenity",
"A4444","Great Mystery Stories","Whodunit","Rodney",5
"A5555","Software Wizardry","Abugov","D.",10

```

For more information, see <http://www.thatjeffsmith.com/archive/2012/05/formatting-query-results-to-csv-in-oracle-sql-developer/>.

1.8.11 Gauges: In the SQL Worksheet and User-Defined Reports

You can use graphical gauges to display query results in the SQL Worksheet and in user-defined reports. In both cases, you need to specify the name of the value column for the gauge data, and minimum and maximum values on the gauge, and the values to be shown as low and high on the gauge (usually between the minimum and maximum values). In the SQL Worksheet, the required structure for the value to be selected is:

```
'SQLDEV:GAUGE:<min>:<max>:<low>:<high>:' || <value-column>
```

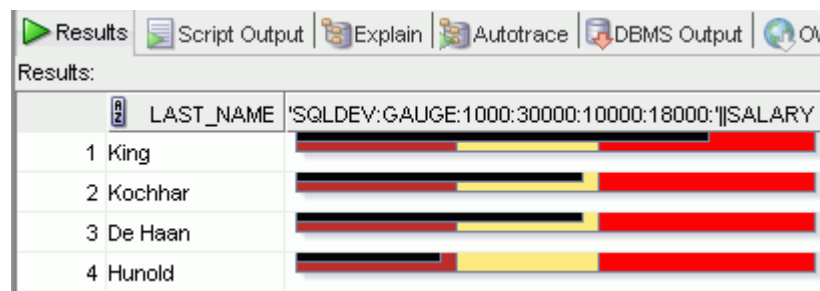
For example, to display the last name and the salary in gauge format, where the gauge shows from 1000 to 30000 with below 10000 as low and above 18000 as high, for employees with ID numbers less than a number to be specified, connect to the supplied HR schema and execute the following query:

```

SELECT last_name, 'SQLDEV:GAUGE:1000:30000:10000:18000:' || salary
       FROM employees WHERE employee_id < :employee_id

```

If you specify 104 as the bind variable value, the output appears as shown in the following figure:



For a user-defined gauge report, the query must specify only the value column, the minimum and maximum values, and the low and high values, and optionally a

WHERE clause. The required structure for the query (before any optional WHERE clause) is:

```
SELECT <value-column>, <min>, <max>, <low>, <high> FROM <table-name>
```

For example, to create a report of salaries in gauge dial format, with the same values and WHERE clause as in the preceding query, right-click on User Defined Reports in the Reports navigator and select **New Report**. In the New Report dialog box, specify a report name; for **Style**, select **Gauge**; and enter the following on the subtabs:

- **SQL Query:**

```
SELECT salary, 1000, 30000, 10000, 18000 FROM employees
WHERE employee_id < :EMPLOYEE_ID;
```

- **Property:**

Gauge Type: Status Meter (or **Dial** if you prefer semicircular dials, each with a label containing the salary amount and a "needle" pointing to an appropriate place on the dial).

Enable (check) **Use Query to Specify the Minimum, Maximum, Low, and High Values**.

Click **Apply**.

Use the Reports navigator to view the newly created user-defined report. For **Connection**, specify one that connects to the HR sample schema. For the bind variable value, specify 104.

Related Topics:

[Using the SQL Worksheet](#) (page 1-41)

[User Defined reports](#) (page 1-89)

1.8.12 Entering OLAP DML Statements in the SQL Worksheet

If Oracle OLAP is available for the connection associated with a SQL Worksheet, you can enter OLAP DML statements by prefixing a tilde (~) to the statement and clicking the Run Statement (for a selected individual statement) or Run Script (for all statements on the worksheet) icon. For example, to enter the `show tod` statement to show the time of day:

```
~show tod
```

To see the output of OLAP DML statements, you must display the DBMS Output Pane (View > Dbms Output).

For information about OLAP DML statements, see the *Oracle OLAP DML Reference* help topic.

1.9 Using Snippets to Insert Code Fragments

Snippets are code fragments, such as SQL functions, Optimizer hints, and miscellaneous PL/SQL programming techniques. Some snippets are just syntax, and others are examples. You can insert and edit snippets when you are using the SQL Worksheet or creating or editing a PL/SQL function or procedure.

To display snippets, from the **View** menu, select **Snippets**. In the snippets window (on the right side), use the drop-down to select a group (such as Aggregate Functions or

Character Functions). In most cases, the fragments in each group do not represent all available objects in that logical grouping, or all formats and options of each fragment shown. For complete and detailed information, see the Oracle Database documentation.

To insert a snippet into your code in a SQL Worksheet or in a PL/SQL function or procedure, drag the snippet from the snippets window and drop it into the desired place in your code; then edit the syntax so that the SQL function is valid in the current context. To see a brief description of a SQL function in a tooltip, hold the pointer over the function name.

For example, you could type `SELECT` and then drag `CONCAT(char1, char2)` from the Character Functions group. Then, edit the `CONCAT` function syntax and type the rest of the statement, such as in the following:

```
SELECT CONCAT(title, ' is a book in the library.') FROM books;
```

Related Topics:

[User-Defined Snippets](#) (page 1-51)

[Running and Debugging Functions and Procedures](#) (page 1-36)

1.9.1 User-Defined Snippets

You can create and edit snippets. User-defined snippets are intended mainly to enable you to supplement the Oracle-supplied snippets, although you are also permitted to replace an Oracle-supplied snippet with your own version.

When you create a user-defined snippet, you can add it to one of the Oracle-supplied snippet categories (such as Aggregate Functions) or to a category that you create. If you add a snippet to an Oracle-supplied category and if your snippet has the same name as an existing snippet, your snippet definition replaces the existing one. (If you later upgrade to a new version of SQL Developer and if you choose to preserve your old settings, your old user-defined snippets will replace any Oracle-supplied snippets of the same name in the new version of SQL Developer.)

To create a snippet, do any of the following:

- Open the Snippets window and click the **Add User Snippets** icon.
- Select text for the snippet in the SQL Worksheet window, right-click, and select **Save Snippet**.
- Click the **Add User Snippet** icon in the Edit Snippets (User-Defined) dialog box.

To edit an existing user-defined snippet, click the **Edit User Snippets** icon in the Snippets window.

Information about user-defined snippets is stored in a file named `UserSnippets.xml` under the directory for user-specific information.

Related Topics:

[Save Snippet \(User-Defined\)](#) (page 6-138)

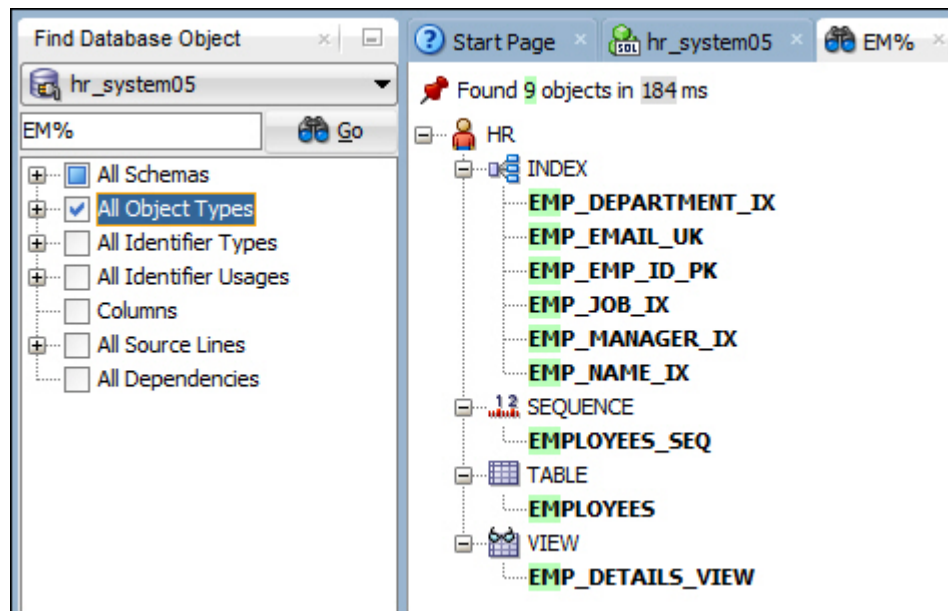
[Edit Snippets \(User-Defined\)](#) (page 6-139)

[Location of User-Related Information](#) (page 1-122)

1.10 Finding Database Objects

You can find various types of objects (tables, columns, declarations within functions or procedures, and so on) associated with an Oracle database connection, and open editing panes to work with those objects. To open the Find Database Object pane, right-click a connection name in the Connections navigator and select **Find DB Object**. You can also click **View**, then **Find DB Object**.

Select the desired database connection, the types of objects for which to search, and a search string. The following figure shows Find Database Objects pane with results from a search for all objects associated with a connection named `hr_system05` where the object name starts with `EM`.



In this example, the search finds nine objects across four object types. In the display of results, you can click any of the object names to bring up the object in an appropriate editor.

Connection: Database connection to use for the search.

Name: An object name or a string containing one or more wildcard characters. For example: `EM%` for all names starting with `EM`.

All Schemas: Select one or more schemas, or all schemas, in which to perform the search. In this example, the blue solid square reflects the fact that the schema you selected is the schema to search; however, you click to select all schemas, or you can expand All Schemas and individually select schemas.

Types: Type of object for which to restrict the search. You can search all columns, or all or individual identifier types, identifier usages, and source lines.

All Dependencies: Extends object search to all dependent objects. For example, if a table is referenced by several views and if that table is found by a search, specifying All Dependencies will also find those dependent views.

Click the **Go** icon to display objects that meet the specified criteria. To view or edit one of the objects (or the parent object that contains the specified object), double-click or right-click its name in the results display.

Related Topics:

[Database Objects](#) (page 1-13)

1.11 Using Recent Objects

You can easily access database objects that you have opened during the current SQL Developer session, even if you have closed them, by using the Recent Objects window. To display this window, click **View**, then **Recent Objects**.

The Recent Objects window displays the names of database objects that you have opened, grouped by object type. You can double-click an object name to open it in the appropriate editor. You can right-click an object name to perform other operations relevant to that object type.

For example, if you were viewing or editing the HR.EMPLOYEES table, then closed the table and disconnected from the HR schema, and if you then double-click EMPLOYEES in the Recent Objects window, the connection is reestablished and the table is opened for editing.

Related Topics:

[Database Objects](#) (page 1-13)

1.12 Using Versioning

SQL Developer provides integrated support for the Subversion and Git versioning and source control systems, and you can add support for other such systems as extensions by clicking Help, then Check for Updates. Available extensions include Subversion and GIT. The SQL Developer documentation does not provide detailed information about the concepts and operations of such systems; it assumes that you know them or can read about them in the product documentation.

- For information about Subversion, see <http://subversion.apache.org/>. For Subversion documentation, see <http://svnbook.red-bean.com/>.
- For information about Git, see <http://git-scm.com/>.

To access the versioning features of SQL Developer, use the **Team** menu.

If you create any versioning repositories, you can use the hierarchical display in the Files navigator, which is marked by a folder icon. (If that navigator is not visible, click **View**, then **Files**.) You can also view a hierarchical display of repositories and their contents in the Versioning navigator.

Related Topics:

[About Subversion and SQL Developer](#) (page 1-53)

[Pending Changes](#) (page 1-54)

[SQL Developer Concepts and Usage](#) (page 1-1)

1.12.1 About Subversion and SQL Developer

Before you can work with a Subversion repository through SQL Developer, you must create a connection to it. When you create a local Subversion repository, a connection to it is automatically created, and this can be seen in the Versioning navigator. You can subsequently edit the connection details.

To connect to an existing Subversion repository (if no connection to that repository exists), use the Versioning navigator, as follows:

1. If the Versioning navigator is not visible, click **View**, then **Team**, then **Versioning Navigator**.
2. Right-click the top-level Subversion node and select **New Repository Connection**.
3. Complete the information in the Subversion: Create/Edit Subversion Connection dialog box, including the URL of the existing repository.
4. Use the connection to access the repository.

Existing files must be imported into the Subversion repository to bring them under version control. Files are then checked out from the Subversion repository to a local folder known as the "Subversion working copy". Files created in (or moved into) SQL Developer must be stored in the Subversion working copy.

Files newly created within SQL Developer must be added to version control. Changed and new files are made available to other users by committing them to the SQL Developer repository. The Subversion working copy can be updated with the contents of the Subversion repository to incorporate changes made by other users.

Related Topics:

[Subversion: Create/Edit Subversion Connection](#) (page 6-143)

1.12.2 Pending Changes

The Pending Changes window is displayed if you request Pending Changes, or when you initiate an action that changes the local source control status of a file. This window shows files that have been added, modified or removed (locally or remotely), files whose content conflicts with other versions of the same file files that have not been added to source control files that are being watched, and files for which editors have been obtained. You can use this information to detect conflicts and to resolve them where possible.

The Outgoing pane shows changes made locally, the Incoming pane shows changes made remotely, and the Candidates pane shows files that have been created locally but not yet added to source control. You can double-click file names to edit them, and you can use the context menu to perform available operations.

1.13 Using DBA Features in SQL Developer

SQL Developer enables users with DBA (database administrator) privileges to view (and in some cases edit) certain information relevant to DBAs and, in some cases, to perform DBA operations. To perform DBA operations, use the **DBA navigator**, which is similar to the Connections navigator in that it has nodes for all defined database connections. (If the DBA navigator is not visible, select **View**, then **DBA**.)

If no connections appear in the DBA navigator display, you must add at least one connection. To add a connection to the DBA navigator, click the plus (+) icon or right-click Connections and select **Add Connection**; then select a connection. (You can also specify a new connection by clicking the plus (+) icon in the Select Connection dialog box.)

You should add only connections for which the associated database user has DBA privileges (or at least privileges for the desired DBA navigator operations) on the

specified database. If you attempt to perform a DBA navigator operation for which the database connection user does not have sufficient privileges, you will receive an error.

The following options are available from the DBA navigator hierarchy:

Connections

<connection-name>

[Container Database \(CDB\)](#) (page 1-57)

[Database Configuration](#) (page 1-57)

[Initialization Parameters](#) (page 1-57)

[Automatic Undo Management](#) (page 1-58)

[Current Database Properties](#) (page 1-58)

[Restore Points](#) (page 1-58)

[View Database Feature Usage](#) (page 1-58)

[Database Status](#) (page 1-58)

[Status](#) (page 1-58)

[DB Instance](#) (page 1-58)

[Data Guard](#) (page 1-59)

[Data Pump](#) (page 1-59)

[Export Jobs](#) (page 1-59)

[Import Jobs](#) (page 1-59)

[Oracle Sharding](#) (page 1-60)

[Performance](#) (page 1-60)

[Snapshots](#) (page 1-61)

[Baselines](#) (page 1-61)

[Baseline Templates](#) (page 1-61)

[Automatic Database Diagnostic Monitor \(ADDM\)](#) (page 1-61)

[ASH \(Active Session History\) Report Viewer](#) (page 1-61)

[AWR \(Automatic Workload Repository\)](#) (page 1-62)

[RMAN Backup/Recovery](#) (page 1-62)

[Backup Jobs](#) (page 1-62)

[Backup Sets](#) (page 1-62)

[Image Copies](#) (page 1-62)

[RMAN Settings](#) (page 1-63)

[Scheduled RMAN Actions](#) (page 1-63)

[Resource Manager](#) (page 1-64)

[Consumer Group Mappings](#) (page 1-64)

[Consumer Groups](#) (page 1-64)

[Plans](#) (page 1-64)

[Settings](#) (page 1-64)

[Statistics](#) (page 1-64)

[SQL Translator FrameWork](#) (page 1-64)

[SQL Translator Profiles](#) (page 1-64)

[SQL Translators](#) (page 1-64)

[Scheduler](#) (page 1-64)

[Global Attributes](#) (page 1-65)

[Job Classes](#) (page 1-65)

[External Destinations](#) (page 1-65)

[Security](#) (page 1-65)

[Audit Settings](#) (page 1-65)

[Profiles](#) (page 1-65)

[Roles](#) (page 1-65)

[Users](#) (page 1-65)

[Storage](#) (page 1-65)

[Archive Logs](#) (page 1-65)

[Control Files](#) (page 1-65)

[Datafiles](#) (page 1-66)

[Redo Log Groups](#) (page 1-66)

[Rollback Segments](#) (page 1-66)

[Tablespaces](#) (page 1-66)

[Temporary Tablespace Groups](#) (page 1-66)

[Tuning](#) (page 1-66)

[Real Time SQL Monitor](#) (page 1-66)

To perform limited database management operations, you can right-click the connection name in the DBA navigator display and select **Manage Database**. For example, if a listener is running with a static listener configured for the database, you can start and stop the database, force database startup, and restrict access to the database.

You can right-click an item (node) at any level in the DBA navigator hierarchy to display a context menu with commands relevant to that item. Typical commands include the following:

- **Refresh** queries the database for the current details about the selected object (for example, a connection or just a table).
- **Apply Filter** restricts the display of objects using a filter that you specify. To remove the effects of applying a filter, right-click the node and select **Clear Filter**.
- **Open** displays a pane with relevant information on the right side of the window. The pane may be read-only or editable.
- **Create New** enables you to create a new object of that type.
- **Help** displays a brief definition or description of objects of that type.

The pane with information about an item opened from the DBA navigator typically contains icons and other controls for the following:

- **Freeze View** (the **pin**) keeps that object's tab and information in the window when you click another object in the Connections navigator; a separate tab and display are created for that other object. If you click the pin again, the object's display is available for reuse.
- **Run** updates the display by querying the database for the latest information.
- **Refresh** lets you select an interval for automatically updating the display.
- **Actions** displays a menu with actions appropriate for the object. The actions are the same as when you right-click an object of that type in the Connections navigator.

With information displays that are in grid form:

- You can often use the context (right-click) menu on a header or a data cell to perform operations like those for the table data grid, as explained in [Entering and Modifying Data](#) (page 1-34). For example, right-clicking on a header lets you select options like Auto-fit, Columns, Sort, Delete Persisted Settings (such as any sort order specifications), and Filter Column; right-clicking any data cell lets you select options like Save Grid as Report, Single-Record View, Count Rows, Find/Highlight (find and optionally highlight values in the grid), and Export (unload data).
- You can click the column heading to display a Filter box to restrict the display to entries containing a string. For example, in the Database Configuration: Current

Database Properties page, you can click PROPERTY_NAME and type NLS to see only properties containing "NLS" (that is, globalization properties).

For information that is read-only (not editable) using the DBA navigator, you may have other options within SQL Developer to specify relevant values. For example, to change the value of any globalization support (NLS) parameter, you can use the Database: NLS preferences pane to change the value for use with all SQL Developer connections (current and future), or you can use the ALTER SESSION statement in the SQL Worksheet window to change the value for the current connection only.

Detailed explanations of various DBA options, including usage and reference information, are available in appropriate manuals in the *Oracle Database Documentation Library* on the Oracle Technology Network (OTN).

1.13.1 Container Database (CDB)

(Available only for Release 12c connections.) Includes options for managing a multitenant container database (CDB) and the pluggable databases (PDBs) within it.

You can right-click to perform the following operations:

- **Create Pluggable Database:** Displays the [Create Pluggable Database](#) (page 6-36) dialog box.
- **Drop Pluggable Database:** Displays the [Drop Pluggable Database](#) (page 6-84) dialog box.
- **Clone Pluggable Database:** Displays the [Clone Pluggable Database](#) (page 6-5) dialog box.
- **Plug In Pluggable Database:** Displays the [Plug In Pluggable Database](#) (page 6-120) dialog box.
- **Unplug Pluggable Database:** Displays the [Unplug Pluggable Database](#) (page 6-156) dialog box.
- **Modify State:** Displays the [Modify Pluggable State](#) (page 6-114) dialog box.
- **Clone PDB to Oracle Cloud:** Displays the [Clone PDB to Oracle Cloud](#) (page 6-4) dialog box.
- **Refresh Pluggable Database:** Refreshes a cloned pluggable database from the source pluggable database, provided the cloned PDB has been opened in read-only mode.
- **Relocate Pluggable Database:** Displays the Relocate Pluggable Database dialog box, which enables you to move a PDB from one container database to another.

1.13.2 Database Configuration

Includes the following options related to database configuration management.

1.13.2.1 Initialization Parameters

For each database initialization parameter, displays the name, value (current value), default value, description, and other information. You can modify the values of some parameters.

1.13.2.2 Automatic Undo Management

Automatic undo management is a mode of the database in which undo data is stored in a dedicated undo tablespace. The only undo management that you must perform is the creation of the undo tablespace; all other undo management is performed automatically. The Automatic Undo Management option displays information about automatic undo management and any recommendations relating to its use. You can change the Retention value.

1.13.2.3 Current Database Properties

For each current database initialization parameter, displays the name, value (current value), and description.

1.13.2.4 Restore Points

Displays restore points that can be used for recovery; lets you create and delete restore points. A **restore point** is a name associated with a timestamp or an SCN of the database. A restore point can be used to flash back a table or the database to the time specified by the restore point without the need to determine the SCN or timestamp. Restore points are also useful in various RMAN operations, including backups and database duplication.

1.13.2.5 View Database Feature Usage

Displays database features and the number of detected usages for each.

1.13.3 Database Status

Includes options for displaying status information about the database.

1.13.3.1 Status

Includes tabs for displaying information about the Database Status, Oracle Host, Oracle Home, and TNS Listener.

1.13.3.2 DB Instance

Displays the database Instance Viewer.

1.13.3.3 Instance Viewer

The database instance viewer enables a DBA user to see a graphical representation of information about the instance associated with a connection. The types of information displayed include:

- Database
- Clients
- Sessions
- Processes (Counts, Execution Rate, Parse Rate, Open Cursors, Commit Rate)
- Waits
- Memory (DB Block Rate, Logical Reads, Allocation, Redo Generation)
- Storage (Files, Redo Log)

- DB CPU Ratio (database operations as a percentage of CPU activity)
- Top SQL (provides performance metrics on SQL operations, which are sorted based on the CPU time required)

You can interact with the instance viewer display, including:

- Panning the diagram by dragging it
- Zooming the diagram in or out by using control-drag
- Drilling down to a detailed display by hovering over graphical element to show a dark blue border and then double-clicking in that element

You can control the amount of history shown in graphs that include a time axis by using the Database: Instance Viewer preferences.

1.13.4 Data Guard

Displays a console for the Oracle Data Guard feature, which is described in detail in *Oracle Data Guard Concepts and Administration*.

When Oracle Data Guard is available for a particular database connection, the Data Guard node is automatically displayed for the connection in the DBA panel.

Console

The Data Gguard console enables a DBA user to see a graphical tree-like representation of the primary and standby databases associated with a connection. You can drill down for more detail by hovering over a graphical element to show a dark blue border and then double-clicking in that element. You can expand or contract a database by clicking the round icon next to it in the console.

You can right-click a database to perform the following operations:

- **Open Instance Viewer:** Opens the database instance viewer using the associated connection.
- **Open Local Values:** Displays information that Data Guard gathers for the database.
- **Edit Connection Association:** Enables you to configure the connection to associate with the database. To find out if a connection already exists, move the cursor over the database to see the hover text.
- **Remove Connection Association:** Removes the associated database connection.

1.13.5 Data Pump

Includes options for using the Oracle Data Pump Export and Import utilities, which are described in detail in *Oracle Database Utilities*.

1.13.5.1 Export Jobs

Displays any Data Pump Export jobs. You can right-click and select **Data Pump Export Wizard** to create a Data Pump Export job.

1.13.5.2 Import Jobs

Displays any Data Pump Import jobs. You can right-click and select **Data Pump Import Wizard** to create a Data Pump Import job.

1.13.6 Oracle Sharding

Includes options for displaying status information about a sharded database.

Before using these options, you should be familiar with the Oracle Sharding and Oracle Data Guard concepts. These features are described in detail in the *Oracle Database Administrator's Guide* and *Oracle Data Guard Concepts and Administration*.

When a shard catalog (an Oracle database that supports shard deployment) is available for a particular database connection, the Oracle Sharding node is automatically displayed for the connection in the DBA panel.

Sharded DB Viewer

The sharded database viewer enables a DBA user to see a graphical representation of information about the sharded database associated with a connection. The viewer displays the regions, shard directors, shard spaces, shard groups, broker sets, and shards (or individual databases) for a sharded database. The red and gray shards signify primary and standby databases respectively. You can drill down for more detail by hovering over a graphical element to show a dark blue border and then double-clicking in that element.

You can right-click a shard to perform the following operations:

- **Open Instance Viewer:** Opens the database instance viewer using the associated connection.
- **Open Data Guard Console:** Opens the Data Guard console if Oracle Data Guard is associated with the database.
- **Edit Connection Association:** Enables you to configure the connection to associate with the shard. To find out if a connection already exists, move the cursor over the shard to see the hover text.
- **Remove Connection Association:** Removes the associated database connection.

1.13.7 Performance

Includes options to create objects and view reports related to the use of database statistics. Oracle Database automatically persists the cumulative and delta values for most of the statistics at all levels (except the session level) in the Automatic Workload Repository (AWR). This process is repeated on a regular time period and the results are captured in an AWR snapshot. The delta values captured by the snapshot represent the changes for each statistic over the time period.

A statistical baseline is a collection of statistic rates usually taken over a time period when the system is performing well at an optimal level. Use statistical baselines to diagnose performance problems by comparing statistics captured in a baseline to those captured during a period of poor performance. This enables you to identify specific statistics that may have increased significantly and could be the cause of the problem. AWR supports the capture of baseline data by enabling you to specify and preserve a pair or range of AWR snapshots as a baseline.

Note:

Most Performance operations are performed internally by SQL Developer through calls to subprograms of the DBMS_WORKLOAD_REPOSITORY package (such as CREATE_SNAPSHOT, CREATE_BASELINE, CREATE_BASELINE_TEMPLATE, and several with names ending in _REPORT_HTML). These are documented in *Oracle Database PL/SQL Packages and Types Reference*.

1.13.7.1 Snapshots

Snapshots are sets of historical data for specific time periods that are used for performance comparisons by Automatic Database Diagnostic Monitor (ADDM). By default, Oracle Database automatically generates snapshots of the performance data once every hour and retains the statistics in AWR for 8 days. You can also manually create snapshots or change the snapshot retention period, but it is usually not necessary.

AWR compares the difference between snapshots to determine which SQL statements to capture based on the effect on the system load. This reduces the number of SQL statements that must be captured over time. After the snapshots are created, ADDM analyzes the data captured in the snapshots to perform its performance analysis.

1.13.7.2 Baselines

A baseline is a set of snapshots from a specific time period that is preserved for comparison with other snapshots when a performance problem occurs. The snapshots contained in a baseline are excluded from the automatic AWR purging process and are retained indefinitely.

1.13.7.3 Baseline Templates

Baseline templates enable you to create baselines for a contiguous time period in the future. There are two types of baseline templates: Single (for a single contiguous time period) and Repeating (based on a repeating time schedule; useful if you want Oracle Database to automatically capture a contiguous time period on an ongoing basis).

1.13.7.4 Automatic Database Diagnostic Monitor (ADDM)

ADDM is an advisor that analyzes Automatic Workload Repository data on a regular basis, diagnoses the root causes of performance problems, provides recommendations for correcting any problems, and identifies non-problem areas of the system.

Related Topics:

Oracle Database Performance Tuning Guide

1.13.7.5 ASH (Active Session History) Report Viewer

ASH enables you to examine and perform detailed analysis on sampled activity of active sessions. To facilitate the identification of transient performance problems, Oracle Database samples active sessions every second. By capturing only active sessions, a manageable set of data is represented with its size being directly related to the work being performed, rather than the number of sessions allowed on the system.

To view the report, specify the desired options (From Time, To Time, Filters) and click the **Generate Report** icon.

Related Topics:

Oracle Database Performance Tuning Guide

1.13.7.6 AWR (Automatic Workload Repository)

An AWR report shows data captured between two snapshots (or two points in time). AWR reports are divided into multiple sections. The content of the report contains the workload profile of the system for the selected range of snapshots. The HTML report includes links that can be used to navigate quickly between sections. The following types of AWR reports are available:

- **AWR Report Viewer:** Generates a Workload Repository Report for a baseline or a snapshot range.
- **Difference Report Viewer:** Generates a Workload Repository Compare Period Report for two baselines or snapshot ranges.
- **SQL Report Viewer:** Generates a Workload Repository SQL Report for a baseline or a snapshot range for a specific SQL statement.

To view a report, specify the desired options and click the **Generate Report** icon.

Related Topics:

Oracle Database Performance Tuning Guide

1.13.8 RMAN Backup/Recovery

Includes options related to database backup and recovery. The options use the Oracle Database Recovery Manager (RMAN) feature, which is described in detail in *Oracle Database Backup and Recovery User's Guide*. You should be familiar with RMAN concepts and techniques before using these options.

In any dialog box or wizard for RMAN operations, you can click the SQL or Summary tab to see the statements that will be used to implement the specified options.

- **Create Backup Wizard** (on context menu) Displays a wizard where you specify Backup Properties (type of backup and whether to use RMAN encryption), and Script Properties (Save to File and where to save the script file, or Run Scheduler Job and specify a credential or create a new one).

For **Create New Credential**, specify the credential name, and the user name and user password on the server system.

1.13.8.1 Backup Jobs

Displays the backup jobs that have been previously run; lets you create and run new backup. (Note that backup jobs are distinct from action jobs.)

1.13.8.2 Backup Sets

Displays the backup sets that have been created by previous backup jobs and that can be used for recovery.

1.13.8.3 Image Copies

Displays the image copies that have been created by previous backup jobs and that can be used for recovery.

1.13.8.4 RMAN Settings

Displays settings for backup and recovery. (These settings are stored in the server and are used and managed by RMAN.)

1.13.8.5 Scheduled RMAN Actions

For Oracle Database 11.1 and later connections: Displays DBMS_SCHEDULER jobs that have been used to execute RMAN scripts; lets you view log files.

1.13.8.6 Using Action Jobs

Action jobs are applicable to connections to an Oracle Database Release 11.1 or later database.

Most backup and recovery actions involve RMAN scripts. The action jobs dialog boxes let you save the generated RMAN to a disk; you can then copy the script to the server system and run the script there.

Before you can execute the script for an action job, you must create a DBMS_SCHEDULER credential by running a procedure in the following format:

```
BEGIN
  DBMS_SCHEDULER.CREATE_CREDENTIAL(
    username => 'user-name',
    password => 'password',
    database_role => NULL,
    windows_domain => NULL,
    comments => NULL,
    credential_name => 'credential-name'
  );
END;
/
```

Then, run the Action Jobs, Configure action. This action places some scripts in the server-side database home (in *<home>/sqldeveloper/dbascripts*) and creates some DBMS_SCHEDULER program objects in the database. (If you want to unconfigure an action job, you can use the Action Jobs, Unconfigure action, which removes the server-side directory containing the script and log files and drops the DBMS_SCHEDULER program objects.)

After you perform the configuration, the Run Scheduler Job action of the Script Processing control becomes available in the RMAN dialog boxes, and you can click **Apply** to cause the RMAN script to be executed in the server using a DBMS_SCHEDULER job.

After an RMAN job has been run, you can view the log file containing the output from RMAN by using the Action Jobs, View Latest Log action. This lets you check for any errors that may have occurred during the running of the RMAN script.

Some RMAN jobs involve performing a database restart. Examples are setting the archive log mode and some whole database backup and restore operations. In such cases, after you click Apply you are asked to confirm that you want to proceed; and if you do proceed, the job is queued (with no waiting for the completion). Because of the restart, the SQL Developer connection must be disconnected and then connected again after the database is restarted. After the reconnection, examine the log file to see whether the job completed successfully.

1.13.9 Resource Manager

Includes the following options related to database resource management.

1.13.9.1 Consumer Group Mappings

A consumer group mapping specifies mapping rules that enable the Resource Manager to automatically assign each session to a consumer group upon session startup, based upon session attributes. The Consumer Group Mappings option displays, for each attribute, its priority, value, and associated consumer group.

1.13.9.2 Consumer Groups

A resource consumer group is a group of sessions that are grouped together based on resource requirements. The Resource Manager allocates resources to resource consumer groups, not to individual sessions. The Consumer Groups option displays, for each consumer group, its description and whether it is mandatory.

1.13.9.3 Plans

A resource plan is a container for directives that specify how resources are allocated to resource consumer group; you specify how the database allocates resources by activating a specific resource plan. The Plans option displays, for each plan, its description and if its status is Active.

1.13.9.4 Settings

Lists any active resource plans.

1.13.9.5 Statistics

Lists various resource-related statistics (if applicable).

1.13.10 SQL Translator Framework

(Available only for Release 12c connections.) Includes options for creating and managing SQL translation profiles and SQL translators.

1.13.10.1 SQL Translator Profiles

A SQL translation profile is a database schema object that directs how non-Oracle SQL statements are translated into Oracle SQL dialects.

1.13.10.2 SQL Translators

A SQL translator is a stored procedure that translates non-Oracle SQL statements into Oracle SQL dialects automatically. After a SQL translation profile is created, users can register a SQL translator with the translation profile to translate the non-Oracle SQL statements.

1.13.11 Scheduler

Includes the following options related to [Scheduling Jobs Using SQL Developer](#) (page 1-67). (The objects under Scheduler in the DBA navigator are for objects that owned by the SYS user and that can be created and modified only by users with DBA privileges. Other objects are listed under Scheduler for users in the Connections navigator.)

1.13.11.1 Global Attributes

The Global Attributes display lets you view and edit attributes such as the default time zone, the email sender and server, event expiry time, log history retention, and maximum job slave processes.

1.13.11.2 Job Classes

The Job Classes display lets you view and edit information about job classes. The information for each job class includes the job class name, logging level, log history, resource consumer group, service, and comments.

1.13.11.3 External Destinations

The External Destinations display lets you view information about external destinations for jobs.

1.13.12 Security

Includes the following options related to database security management. For profiles, roles, and users, you can perform relevant operations, such as creating new objects of that type and editing and dropping existing objects.

1.13.12.1 Audit Settings

The Audit Settings display includes the audit trail setting, whether SYS user operations are audited, and the directory or folder for the audit file.

1.13.12.2 Profiles

A profile is a set of limits on database resources. If you assign the profile to a user, then that user cannot exceed these limits. The Profiles option displays any limits on activities and resource usage for each profile.

1.13.12.3 Roles

A role is a set of privileges that can be granted to users or to other roles; you can use roles to administer database privileges. The Roles option displays the roles and their authentication settings.

1.13.12.4 Users

A database user is an account through which you can connect to the database. The Users option displays status and usage information about each database user.

1.13.13 Storage

Includes the following options related to database storage management.

1.13.13.1 Archive Logs

A database user is an account through which you can connect to the database. The Users option displays status and usage information about each database user.

1.13.13.2 Control Files

A control file is a binary file that records the physical structure of a database and contains the names and locations of redo log files, the time stamp of the database creation, the current log sequence number, checkpoint information, and so on. The

Control Files option displays, for each control file, its status, file name, and file directory.

1.13.13.3 Datafiles

A data file is a physical file on disk that was created by Oracle Database and contains the data for a database. The data files can be located either in an operating system file system or Oracle ASM disk group. The Datafiles option displays, for each data file, its file name, tablespace, status, and other information.

1.13.13.4 Redo Log Groups

A redo log group contains one or more members: each online redo log member (which corresponds to an online redo log file) belongs to a redo log group. The contents of all members of a redo log group are identical. The Redo Log Groups option displays, for each redo log group, its status, number of members, and other information.

1.13.13.5 Rollback Segments

A rollback segment records the before-images of changes to the database. The Rollback Segments option displays, for each rollback segment, its name, status, tablespace, and other information.

1.13.13.6 Tablespaces

A tablespace is a database storage unit that groups related logical structures together. The database data files are stored in tablespaces. The Tablespaces option displays, for each tablespace, its name; megabytes allocated, free, and used; and other information.

1.13.13.7 Temporary Tablespace Groups

A temporary tablespace group is a tablespace group that is assigned as the default temporary tablespace for the database. (A tablespace group enables a database user to consume temporary space from multiple tablespaces. Using a tablespace group, rather than a single temporary tablespace, can alleviate problems caused where one tablespace is inadequate to hold the results of a sort, particularly on a table that has many partitions.) The Temporary Tablespace Groups option displays, for each tablespace group, its name, the number of tablespaces in the group, the total size of the tablespaces, and whether the group is the default temporary tablespace.

1.13.14 Tuning

Provides an option to monitor the database performance in real time.

1.13.14.1 Real Time SQL Monitor

The real time SQL monitor (Oracle Database Release 11.1 and later only) requires the Oracle Tuning Pack. This feature enables you to monitor the performance of SQL statements while they are executing. It displays information about any query currently executing and queries that are done executing for a selected connection. To see detailed information about a query, click the entry. The information is especially useful for real-time monitoring of long-running SQL statements. Cursor statistics (such as CPU times and IO times) and execution plan statistics (such as number of output rows, memory, and temporary space used) are updated close to real-time during statement execution. (Internally, this feature calls the `DBMS_SQLTUNE.REPORT_SQL_MONITOR` subprogram.)

1.14 Scheduling Jobs Using SQL Developer

SQL Developer provides a graphical interface for using the DBMS_SCHEDULER PL/SQL package to work with Oracle Scheduler objects. To use the SQL Developer scheduling features, you must first understand the concepts and essential tasks for job scheduling, which are explained in the chapters about Oracle Scheduler concepts and scheduling jobs in *Oracle Database Administrator's Guide*.

The Scheduler node for a connection appears in the Connections navigator and in the DBA navigator. The DBA navigator displays Scheduler objects owned by the SYS user and requiring DBA privileges for creating and modifying; the Connections navigator displays other Scheduler objects that may require specific privileges for creating and modifying. The Oracle Scheduler objects are grouped under object types that include:

- [Chains](#) (page 1-14)
- [Credentials](#) (page 1-15)
- [Database Destinations](#) (page 1-15)
- [File Watchers](#) (page 1-16)
- [Jobs](#) (page 1-18)
- [Job Classes](#) (page 1-18)
- [Programs](#) (page 1-21)
- [Schedules](#) (page 1-22)
- [Window Groups](#) (page 1-27)
- [Windows](#) (page 1-27)

You can right-click an item (node) at any level in the Scheduler hierarchy to display a context menu with commands relevant to that item. Typical commands include the following:

- **New** [*object-type*] enables you to create a new object of that type.
- **Refresh** queries the database for the current details about the selected object (for example, a connection or just a table).
- **Apply Filter** restricts the display of objects using a filter that you specify. To remove the effects of applying a filter, right-click the node and select **Clear Filter**.
- **Open** displays a pane with relevant information on the right side of the window. The pane may be read-only or editable.
- **Help** displays a brief definition or description of objects of that type.
- **Drop** (with an object selected) drops (deletes) the object.
- **Create Like** (with an object selected) creates an object of that type using information from the selected object. You must specify a name for the newly created object, and you can change other properties as needed.

If you right-click the top-level Scheduler node, you can create an object of a specified type, and you can display the Scheduler Design Editor.

The pane with information about an item opened from the Scheduler hierarchy typically contains icons and other controls for the following:

- **Freeze View** (the **pin**) keeps that object's tab and information in the window when you click another object in the Connections navigator; a separate tab and display are created for that other object. If you click the pin again, the object's display is available for reuse.
- **Run** updates the display by querying the database for the latest information.
- **Refresh** lets you select an interval for automatically updating the display.
- **Actions** displays a menu with actions appropriate for the object. For example, the actions for a job include New Job, Drop, Run Job Now, and Create Like.

With information displays that are in grid form, you can often use the context (right-click) menu on a header or a data cell to perform operations like those for the table data grid. For example, right-clicking on a header lets you select options like Auto-fit, Columns, Sort, Delete Persisted Settings (such as any sort order specifications), and Filter Column; right-clicking any data cell lets you select options like Save Grid as Report, Single-Record View, Count Rows, Find/Highlight (find and optionally highlight values in the grid), and Export (unload data).

Several Scheduler reports are also available.

1.14.1 Scheduler Design Editor

The scheduler design editor is a graphical interface for scheduling jobs using SQL Developer. To create, edit, or delete any Oracle Scheduler objects in the editor, you must be in **write mode**. You can use the editor toolbar toggle button or the right-click command to *Switch to write mode* and *Switch to read-only mode*.

To copy existing objects to the schedule design editor, drag the objects from the Connections navigator onto the editing canvas.

To create a new object, right-click in the canvas and select the context menu command for creating a new object of the desired type.

To open an object, double-click it or right-click it and select **Open**.

To delete an object, right-click it and select **Drop**. You will be prompted to confirm that you want to drop the object.

To remove the object from the editing canvas without deleting the object itself, right-click it and select **Clear**.

To enable a disabled object, right-click it and select **Enable**; to disable an enabled object, right-click it and select **Disable**.

You can click, drag, and release to *set* a relationship between compatible objects. For example, to set a schedule for a job, you can drag the schedule onto the job, which causes a line to be drawn between the objects. You can *unset* (remove) that relationship by right-clicking the line and selecting the appropriate command (such as **Unset Schedule**), and then clicking **OK** when you are prompted to confirm the action.

1.15 Using the Cart

The Cart is a convenient tool for collecting Oracle database objects from one or more database connections, and deploying, exporting, comparing, or copying those objects. You can put objects into one or more carts, each with its own tab. When the Cart

window is opened, it contains an empty cart, although you can create new carts and open previously saved carts in new or existing cart tabs.

For each of the Cart database utilities (Export, Diff [compare], Copy), you can create, save, and open utility-specific configuration settings.

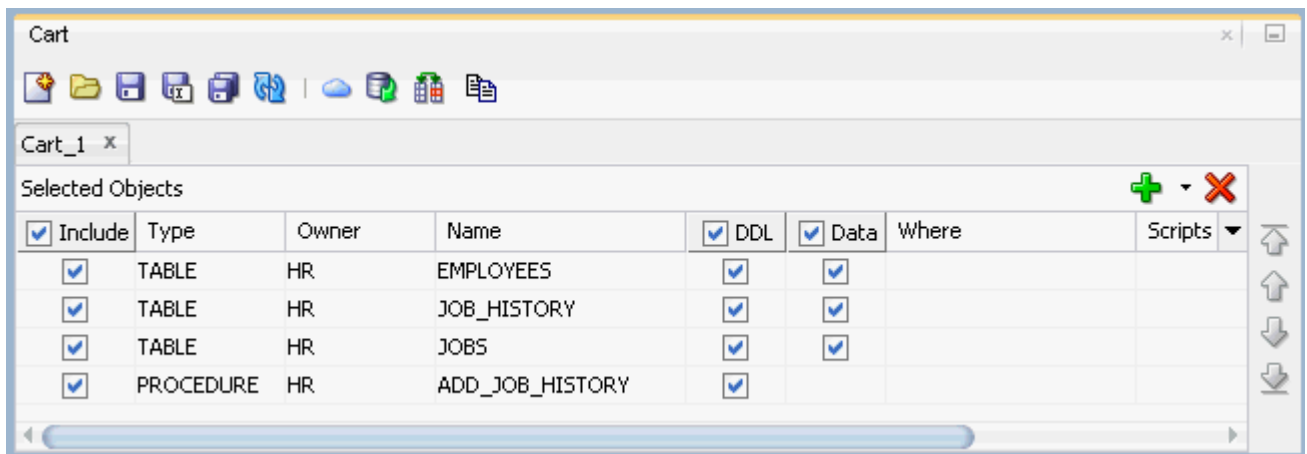
Objects in the Cart are not automatically synchronized with database activity; to update the contents of the cart with the current state of the database, click the Refresh icon. If an object does not exist after a refresh, the object is disabled in the Cart and is underlined to indicate the error.

You can put objects into a cart tab in several ways:

- Drag and drop objects from the Connections navigator into the Cart window.
- Select one or more objects in the Connections navigator, right-click, and select **Add to Cart**.
- Open a previously saved Cart XML file.
- Add scripts By using the Scripts icon drop-down (Add Initial Script, Add Final Script). (If you use a Cart tool that does not support scripts, they are ignored.)

Before you use the Cart, ensure that the Database: Cart and Database: Cart: Cart Deploy preferences reflect what you want.

To display the Cart window, click **View**, then **Cart**. The following figure shows that window after some objects have been added in the default cart tab.



New Cart icon: Opens a new empty cart tab. For this and any other card tab, you can right-click the tab name and select from the following:

- **Close:** Closes the current cart tab.
- **Close Others:** Closes all cart tabs except the current one.
- **Close All:** Closes all cart tabs.
- **Rename:** Renames the current cart tab (for example, if you wanted to change Cart_1 to HR_objects).

Open Cart icon: Displays a dialog box where you specify:

- **Open Cart:** Location and name of the XML file for a previously saved cart. (If the cart file does not exist when you click **Apply**, the Open Cart Error box is

displayed. Ensure that you specify the correct location and file name, and try again.)

- **Open in:** Whether to open the XML file in a new cart tab or in the currently selected cart tab. If in the selected cart, specify whether you want to remove the current objects from the cart before opening the saved cart or to add the objects in the cart file to the current cart.

Save Cart, Save As Cart, Save All Carts icons (active if one or more objects are selected): Displays a dialog box for saving the current cart or all carts to an XML file that you can later open. Specify:

- **File:** Location and name of the XML file in which to save the cart.
- **Encoding:** Character set to be used for encoding of the output file.

For Save All Carts, the dialog box is displayed repeatedly, once for each currently open cart tab.

Refresh Cart (active if one or more objects are selected): Refreshes the current cart tab; validates that what is in the cart is actually in the database.

Deploy Cloud (Deploy to the Oracle Cloud; active if one or more objects are included in the selected cart): After you have added the objects to be deployed and specified any desired options, click this icon to display the Deploy Objects to Cloud dialog box, where you can specify further options and deploy the objects to the Oracle Cloud.

Export (Create a deployment file; active if one or more objects are included in the selected cart): After you have added the objects to be deployed and specified any desired options, click this icon to display the Export Objects dialog box, where you can specify further options and generate the deployment scripts and optionally a .zip file that contains them.

Database Diff: Displays the Diff Objects dialog box, for comparing the selected objects with the objects in another currently open cart tab or a database connection that has access to the destination objects to be compared.

Copy: Displays the Copy Objects dialog box, for copying the selected objects to a database connection.

Selected Objects: Contains objects that have been selected for inclusion in each open cart.

Delete Row icon (red X): Lets you delete selected objects from the Selected Objects area.

Include: Lets you select and deselect all objects or selected objects for inclusion in a Cart operation.

Type: Type of database object.

Owner: Owner of the database object.

Name: Name of the database object.

DDL: Lets you include or exclude the DDL (Data Definition Language) statements for a Cart operation.

Data: For tables and views, lets you include or exclude the associated data for a Cart operation.

Where: Optional WHERE clause (without the WHERE keyword) to filter the data rows from a table or view.

Scripts (applies to Deploy to Cloud, Export, and Copy): Lets you optionally include one or more .sql scripts to be executed at specified locations in the generated master script. For each object, you can click the pencil icon in the Scripts cell; this displays a dialog box to specify the file and character set encoding for any combination of the following:

- **Before Create:** Script to run before the object is created.
- **After Create:** Script to run after the object is created.
- **Before Load** (if data is included): Script to run before the data is loaded into the created object.
- **After Load** (if data is included): Script to run after the data is loaded into the created object.

The default location for these scripts is the Open Script Directory preference under the Database: Utilities: Cart user preferences.

Connection: Connection from which the object was added. You can click the pencil icon in a Connection cell to display the Map Connection dialog box, where you can change the Connection value for one or more objects in the currently selected cart.

Arrow buttons: Let you move objects up or down in the order of Selected Objects. This affects the order in which objects are created in the generated scripts.

1.15.1 Cart Versus Database Export Wizard

The Cart and the Database Export (Unload Database Objects and Data) wizard enable you to perform the same basic operations. A suggested guideline for choosing which one to use is *repeatability* of the deployment/export operation:

- If you want to be able to repeat the operation (with or without slight modifications), the Cart provides a convenient way to save a cart and to deploy it, and then run the scripts in multiple destinations (schemas in the same database or another database). You can also open a saved cart, modify it as needed, save it to a different cart name, and deploy that cart.
- If you will not need to repeat a specific export operation, the wizard provides a convenient interface for such one-time actions.

Related Topics:

[Database Export \(Unload Database Objects and Data\)](#) (page 6-77)

1.16 Spatial Support in SQL Developer

SQL Developer lets you work with Oracle Spatial and Graph data and metadata. Oracle Spatial and Graph (referred to as Spatial and Graph) facilitates the storage, retrieval, update, and query of collections of Spatial and Graph features in an Oracle database. Spatial and Graph includes a schema (MDSYS), a data type (SDO_GEOMETRY), a spatial indexing mechanism and index type, and operators, functions, and procedures for performing area-of-interest queries, spatial join queries, and other spatial analysis operations.

SQL Developer support for Spatial and Graph includes the following:

- [Context Menu Operations on Spatial Data and Metadata](#) (page 1-72)

- [Map Visualization of Spatial Data](#) (page 1-72)

1.16.1 Context Menu Operations on Spatial Data and Metadata

In the Connections navigator display, the context menu (right-click) for a table includes a **Spatial** submenu, which includes the following commands:


- **Update Spatial Metadata** displays the Maintain Spatial Metadata dialog box, which enables you to update the spatial metadata for data in columns of type SDO_GEOMETRY in the table.
- **Validate Geometry Using Tolerance** performs a consistency check, based on the tolerance value that you specify, for valid geometry types in the spatial data in the table, and returns context information if any geometries are invalid. Displays the Validate Geometry dialog box.
- **Validate Geometry Using Dimension Information** performs a consistency check, based on the dimensional information associated with the spatial data (using the USER_SDO_GEOM_METADATA view), for valid geometry types in the spatial data in the table, and returns context information if any geometries are invalid. Displays the Validate Geometry dialog box.
- **Drop Spatial Metadata** deletes the spatial metadata for the specified spatial column (but does not delete the column definition or the data).
- **Create Spatial Index** creates a spatial index (INDEXTYPE IS MDSYS.SPATIAL_INDEX) on the specified spatial geometry column.
- **Drop Spatial Index** deletes the specified spatial index.

1.16.2 Map Visualization of Spatial Data

You can use the Map View window for displaying spatial geometry objects that are returned by a query. The SQL Developer Map View display lacks the rich features of a visualizing tool such as Oracle MapViewer, but it provides a quick and flexible way to see simple representations of spatial data.

To display the Map View, click **View**, then **Map View**. You can resize and reposition the Map View window so that it adequately displays the geometries you want to see.

The following figure show the Map View window display of a query that returns all geometries for counties in the U.S. state of South Carolina. On the left is a map display pane with some icons above it. On the right is a Query List pane with one query (named `South Carolina Counties`) and some icons above it. (Not shown is the database connection drop-down selector to the right of the Query List.)



You can visualize geometries in the Map View window by creating and executing a SQL query, or by selecting a geometry in a table data grid cell and using the context menu:

- [Visualizing Geometries by Creating and Executing a Query](#) (page 1-73)
- [Visualizing Geometries from the Table Data Grid](#) (page 1-74)

1.16.2.1 Visualizing Geometries by Creating and Executing a Query

To visualize geometries by creating and executing a SQL query, follow these steps:

1. Click **View**, then **Map View** to display the Map View window.
2. Optionally, resize and reposition the Map View window.
3. Using the connection selector drop-down on the right above Query List (not shown in the figure), select the database connection to be used for the query.
4. Click the Add New Query (+) icon above Query List to display a query dialog box, and specify the information for the query:

Map Title: Short descriptive name for the display to result from the query.

Example: South Carolina Counties

Map Query: SELECT statement to return the desired geometries. Example (for a table named COUNTIES that includes a geometry column named GEOM and a STATE_ABRV column for 2-character state abbreviations):
`SELECT geom FROM counties WHERE state_abrv = 'SC'`

Map Styles: Optionally, customize the **Line Color** and **Fill Color** for the geometries, and select an **Annotation Column**. For the colors, you can click the square to display a dialog box for specifying a color using a swatch or an HSB or RGB value.

5. Click **OK** to complete the definition.

6. Optionally, repeat steps 4 and 5 to create one or more additional queries. (For a query using a different database connection, repeat steps 3 through 5.)

To execute the query and display the map in the left pane of the window, select (check mark) the query and click the Execute Checked Queries icon above Query List.

To modify a query, select (check mark) it and click the Edit Query (pencil) icon above Query List.

To delete a query from the list, select (check mark) it and click the Remove Query (X) icon above Query List.

In the map display pane on the left, you can click icons above the pane to control aspects of the display:

- Zoom In expands the display and makes the objects appear larger.
- Zoom Out contracts the display and makes the objects appear smaller.
- Fit adjusts the display so that all objects fit in the pane.
- Marquee Zoom lets you select a rectangle (press, drag, release), after which the display zooms to fit the selected rectangle in the pane.

1.16.2.2 Visualizing Geometries from the Table Data Grid

To visualize one or more geometries from the table data grid, follow these steps:

1. Display the information for a spatial table by clicking its name in the Connections navigator hierarchy display.
2. Click the **Data** tab to display the data in the table.
3. In the grid display of table data, click the cell with the spatial data for a geometry object. For example, if the spatial column is named `GEOM`, click one of the cells under `GEOM`. (The cell content will start with `MDSYS . SDO_GEOMETRY`.)
4. Right-click, and select one of the following from the context menu:
 - **Display Geometry Shape** displays the selected geometry object in a standalone box.
 - **Identify Geometry Shape in Map View** highlights the selected geometry object in the Map View window (if it is included in the window).
 - **Invoke Map View on Result Set** displays all geometry objects in the column (that is, it displays the layer).

1.17 Change Manager Support in SQL Developer

A **change plan** is an Oracle Change Manager container for schema object changes. Using SQL Developer, you can create change plans and populate them change information. Later, a database administrator deploys the change plan to one or more destination databases using Oracle Enterprise Manager.

A change plan can contain multiple **change items**, each of which describes changes to be made to a single schema object. There are three types of change item, reflecting three basic actions to be carried out at the destination database:

- An **Add change item** adds a database object.

- A **Drop change item** drops the database object. (The definition of the object to be dropped is copied into the change item, enabling Change Manager to check the definition in the destination database to ensure that the object definitions are the same.)
- A **Modify change item** makes one or more changes to a database object (for example, adding a column or a constraint to a table).

To deploy a change plan to a destination database, the DBA uses Enterprise Manager to create a Change Manager Schema Synchronization from the contents of the change plan. During this process, Change Manager detects and reports conflicts between the change plan and the database. (For example, the database may already contain an object that the change plan is trying to add.) The generated script carries out operations in the correct order.

The Change Management Parameters pane contains user preferences that affect Change Manager support in SQL Developer, especially how objects are compared to determine if a change has occurred and how extensive the DDL statements are when the script for deploying changes is generated.

See the following subtopics for more information related to support for Change Manager:

- [Change Plan Administrators and Developers](#) (page 1-75)
- [Developing Change Plans with SQL Developer](#) (page 1-76)
- [Change Manager Actions to Perform with Enterprise Manager](#) (page 1-78)

Related Topics:

[Change Management Parameters](#) (page 1-93)

1.17.1 Change Plan Administrators and Developers

Change plans are stored in the Enterprise Manager repository. An Enterprise Manager super administrator has unlimited access to all change plans. However, in most cases, it is preferable to create Enterprise Manager administrators with specific Change Plan privileges and limited access to other Enterprise Manager facilities. There are two general types of Change Plan privilege:

- A **Change Plan Administrator** has unlimited access to change plans: the administrator can create, edit, and delete any change plan. A Change Plan Administrator is typically a DBA.
- A **Change Plan Developer** has access to specific change plans. The access may be view-only, or view and edit. A Change Plan Developer is typically a developer of database applications.

To create a Change Plan Administrator or one or more Change Plan Developers, you must use Enterprise Manager.

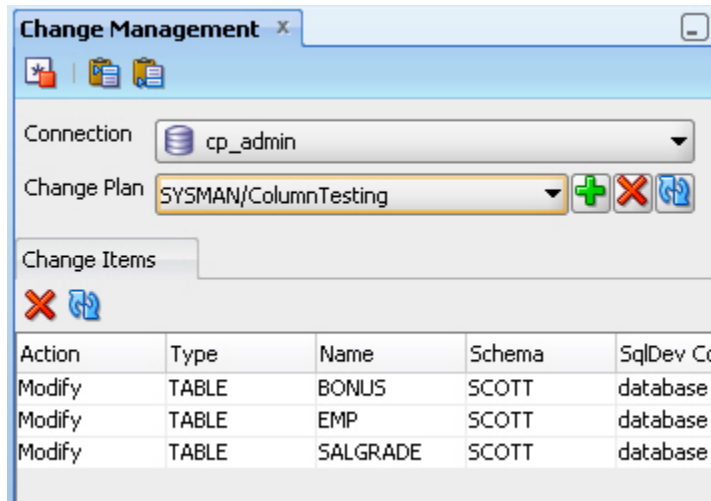
Related Topics:

[Change Manager Actions to Perform with Enterprise Manager](#) (page 1-78)

1.17.2 Developing Change Plans with SQL Developer

The SQL Developer user interface support for Change Manager includes the Change Management window and the Change Management Parameters user preferences.

To display the Change Management window, click View, then Change Management. The Change Management window is displayed within the SQL Developer main window.



At the top of the Change Management window are icons for the following operations:

- Disconnect disconnects the selected Change Plan connection.
- Pre Script lets you enter a script with statements to be executed before the selected change plan is deployed.
- Post Script lets you enter a script with statements to be executed after the selected change plan is deployed.

Connection: Database connection for the Change Management repository.

Change Plan: A change plan within the selected repository. To the right of the Change Plan are icons for the following operations: **Add Plan (+)** to create a new change plan to the repository, **Remove Plan (X)** to delete the selected plan from the repository, and **Refresh** to refresh the window with the current contents of the repository.

Change Items: A list of change items for the selected change plan, with the following information about each item: Action, Database Object Type, Object Name, Schema, SQL Developer Connection, Enterprise Manager Connection, Global Name, Description.

To delete a change item, select its row and click the **Remove Item(s) (X)** icon. To update the change items display to reflect the contents of the repository, click the **Refresh** icon.

Related Topics:

[Creating a Repository Connection](#) (page 1-77)

[Creating and Deleting Change Plans](#) (page 1-77)

[Adding and Updating Change Items](#) (page 1-77)

[Using Change Plans to Capture Change in SQL Developer](#) (page 1-77)

1.17.2.1 Creating a Repository Connection

To store and retrieve change plans, you must create a SQL Developer database connection to the Enterprise Manager repository database. This connection typically logs in as one of the repository database users created as a Change Plan administrator or developer. Create the database connection specifying the user name and password of a suitably privileged user, and the connection details for the schema Enterprise Manager repository.

1.17.2.2 Creating and Deleting Change Plans

To create a change plan:

1. If the Change Management window is not visible, Click **View**, then **Change Management** to display the window.
2. For **Connection**, select the database connection for the Enterprise Manager repository.
3. To the right of Change Plan, click the **Add Plan (+)** icon.
4. In the Create Plan dialog box, specify a name for the plan, and click **Apply**.

To delete a change plan.

1. For **Connection**, select the database connection for the Enterprise Manager repository.
2. For **Change Plan**, select the change plan to be deleted.
3. To the right of Change Plan, click the **Remove Plan (X)** icon.
4. In the Delete Plan dialog box, click **Apply** to confirm that you want to delete the change plan.

1.17.2.3 Adding and Updating Change Items

To add a change item to a change plan or to update (refresh) the change items in a change plan, you must be connected to the Enterprise repository as a user that has Edit privileges on the plan (or that has the Manage Change Plans resource privilege). In the Change Management window, select the relevant change plan from the Change Plan list. The current contents of the plan (if any) are listed in the Change Items display.

To add a change item to the selected change plan, drag a database object from the Connections navigator (for example, a table named EMPLOYEES from a connection named HR) into the Change Items pane. In the dialog box, specify the type of change item to be created: **Add Change Item**, **Drop Change Item**, or **Modify Change Item**.

1.17.2.4 Using Change Plans to Capture Change in SQL Developer

There are two basic approaches to capturing changes in change plans: single database or multiple databases. The approach chosen may depend on the complexity of the work and the number of developers involved.

- **Single database:** Before development starts, the developer creates Monitor change items for all objects that might be modified or dropped. The developer then begins development work in the database -- adding, modifying, and dropping objects. As this work progresses, the developer updates Monitor change items to

make them into Modify change items, or converts them to Drop change items, and creates new Add change items. At the end of development work, the change plan contains the changes that took place during the development cycle.

- **Multiple databases:** One database represents the pre-development state and is not modified during development work. All work takes place in one or more development databases. Monitor and Drop change items are created as needed from the pre-development database, and then updated from the development database(s). Changes from multiple developers working in multiple databases can be combined into a single change plan.

1.17.3 Change Manager Actions to Perform with Enterprise Manager

You must use Enterprise Manager to perform any of the following operations related to change plans:

- Creating change plan administrators and developers
- Creating change plans from Change Manager schema comparisons
- Deploying change plans

For information about performing these operations, see the Enterprise Manager help and documentation.

1.17.3.1 Creating Change Plans from Change Manager Schema Comparisons

You can create change plans in Enterprise Manager and populate them from Change Manager schema comparisons. On the Databases page, select **Schema Change Plans**, then **Create**. Enter the name of the new change plan and click **OK**. In the resulting Change Plan page, click **Create from Comparison**. Select a version of a Schema Comparison, and a Conversion Assignment as follows:

- For the Change From side, select the side of the comparison that represents the original state of the schema objects.
- For the Change To side, select the side of the comparison that represents the developed (or evolved) state of the schema objects.

Based on your selection, Change Manager populates the change plan with Add, Drop, and Modify change items that, if applied to the Change From objects, will convert them into the Change To objects.

1.17.3.2 Deploying Change Plans

To deploy change plans, Change Manager uses the script generation capability contained in its Schema Synchronizations application. You deploy a change plan by creating a synchronization that applies the changes in the change plan to a database that you select. To reduce the chance of incompatible change conflicts, the objects in the database should be identical or similar to the initial objects from which the change plan's change items were created.

To deploy a change plan to a database, select the plan on the Change Plans page, and then click **Create Synchronization from Change Plan**. This takes you to the first page of the Create Schema Synchronization wizard. The process of creating a synchronization from a Change Plan is similar to creating a synchronization from scratch. However, in the first phase of synchronization processing, rather than comparing the source database (or baseline version) to the destination database, Change Manager applies the changes contained in the change plan to the destination

database definitions (not to the actual objects). The result of this operation provides the basis for script generation and execution, the second and third phases of schema synchronization processing.

Conflicts

During change plan processing, Change Manager may detect conflicts between change items and the destination database. In general, a conflict means that the state of the destination database is not what was expected. Examples of conflicts include:

- A table that is to be added already exists, and has a different definition than that contained in the change item.
- A table to be modified does not exist.
- A column to be added to a table already exists and has a different data type than the column added by the change item.
- A constraint to be dropped from a table does not exist.
- A view to be dropped exists, but has a different definition than that stored in the change item.
- A procedure to be dropped does not exist.

After changes have been applied to the destination definitions, Change Manager reports the results of the operations, including any conflicts encountered.

Change Plans and Schema Synchronizations: Comparison

Schema change plans and schema synchronizations are similar in some ways, and use some of the same underlying technology. However, they have different purposes and uses:

- The purpose of a *schema synchronization* is to make the set of objects at the destination database identical to those at the source database or baseline version. Schema synchronization compares the objects at the source and destination, and then generates a script to reconcile all the differences.
- The purpose of a *change plan* is to carry out a specific set of changes at any destination database. The changes may apply to an entire object (Add or Drop), or to attributes of an object (Modify, for example, add a column).

In other words, schema synchronization picks up any changes that have been made to the destination object or schema and makes those changes at the destination. Change plan deployment, however, makes only those changes specified in the change items.

With schema synchronizations, there is no chance of conflict because the source definition always overrides the destination definition. With change plans, however, there is a possibility of conflict, because the specified changes might not be applicable at a destination database or might conflict with other changes.

1.18 Apache Hadoop Connectors Support in SQL Developer

SQL Developer includes support for using connectors for Apache Hadoop to:

- Copy data from an Oracle Database table to the Apache Hadoop Distributed File System (HDFS) and create an Apache Hive external table over that data using **Copy to Hadoop** (a feature of Oracle Big Data SQL)

- Load data quickly from Hadoop to a table in Oracle Database, using **Oracle Loader for Hadoop**
- Load data from Hadoop into an Oracle Database table using SQL and external tables, and query data in-place via external tables, using **Oracle SQL Connector for HDFS**

To use the Apache Hadoop connectors, you must do the following in SQL Developer:

- Click **Tools > Features > Oracle SQL Developer — Hadoop Connectors**
- Right-click each Oracle Database connection that you plan to use, and select **Manage Features > Hadoop Integration** (a toggle setting).

The features require the relevant product to be installed in the Apache Hadoop cluster and Oracle Database. The features are supported for any Hadoop cluster and Oracle Database supported by the products that they use.

The features require JDBC drivers for Apache Hive. To add the JDBC drivers, click **Tools > Preferences > Database > Third Party JDBC drivers**.

Related Topics:

[Database: Third Party JDBC Drivers](#) (page 1-99)

1.18.1 Copy to Hadoop: Overview and Requirements

Copy to Hadoop is a feature of Oracle Big Data SQL. SQL Developer makes it easy to identify and copy data from an Oracle Database table to the Apache Hadoop Distributed File System (HDFS) and create an Apache Hive external table over that data. The data that is copied can be stored as Oracle Data Pump files, Parquet files, or Optimized Row Columnar (ORC) format files. When stored as Oracle Data Pump files the data is retained as Oracle data types, providing an exact copy of the data in the database without any errors due to data type conversions. The data that is copied can be queried by Hive (and Hadoop technologies that access data through Hive), and by Oracle Big Data SQL.

Using SQL Developer, you can copy data and create a new Hive table, or append data to an existing Hive external table that was created by Copy to Hadoop.

Copy to Hadoop is supported with:

- Any Oracle Database that is supported by Oracle Big Data SQL, and with Oracle Database release 11.2 and later.
- Any Hadoop version that is supported by Oracle Big Data SQL.

You can initiate Copy to Hadoop in SQL Developer by right-clicking the Tables icon under any Hive schema. You can initiate Append to an Existing Hive External Table by right-clicking the icon for that Hive table.

Verification of Copy to Hadoop Installation

Copy to Hadoop is installed as part of Oracle Big Data SQL installation. Using SQL Developer with Copy to Hadoop requires the following, which are created during Oracle Big Data SQL installation:

- Database directory object `ORACLE_BIGDATASQL_CONFIG`, which points to the installation directory for Oracle Big Data SQL on the database system
- Database directory object `ORA_BIGDATA_CL_<hadoop_cluster_name>`

- File `hadoop_<hadoop_cluster_name>.env` in the parent directory of the file system location pointed to by `ORACLE_BIGDATASQL_CONFIG`

Check your Big Data SQL installation to verify that these directory objects and file exist.

Setup Required to Use SQL Developer with Copy to Hadoop

Using SQL Developer to run a Copy to Hadoop job with the Direct Copy method requires an SSH connection to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. The SSH connection is selected in the wizard. The file `.sqldev_cp2hadoop_env` must exist in the home directory of the user of that SSH connection. The following is a sample `.sqldev_cp2hadoop_env` file:

```
#!/bin/bash
export CP2HADOOP_HOME=<location of the Copy to Hadoop jars. For example,/opt/oracle/orahivedp>
export HADOOP_CLASSPATH=${CP2HADOOP_HOME}/jlib/*
export WALLET_LOCATION=<location of the wallet files, if using Oracle Wallet. For example,/home/${USER}/wallet>
export TNS_ADMIN=<location of the wallet files, if using Oracle Wallet. TNS_ADMIN should point to the same location. For example, /home/${USER}/wallet>
export CLUSTER_WALLET_LOCATION=${WALLET_LOCATION}
export CLUSTER_TNS_ADMIN=${TNS_ADMIN}
```

Additional setup steps are required on the Hadoop client, the Hadoop edge node, or the node in the Hadoop cluster:

- Create a directory `/home/<SSH user name>/sqldev_temp`. This is a temporary directory used by SQL Developer to store deployment scripts.
- Create the HDFS root directory for storing data files copied from Oracle Database. Typically, this is `/user/<SSH user name>/oracle_warehouse`. Log in to the Hadoop node as a SSH user, and enter:

```
$ hadoop fs -mkdir -p /user/<SSH user name>/oracle_warehouse
```

(Note: Additionally, create the HDFS directory `/user/oracle/oracle_warehouse` if using Copy to Hadoop with Stage method.)

1.18.2 Oracle Loader for Hadoop: Overview and Requirements

Oracle Loader for Hadoop is a tool for high-speed loading of data from Hadoop into a table in Oracle Database. It loads data in parallel and can use direct path load when the database table is partitioned. Oracle Loader for Hadoop runs as a MapReduce job. It partitions and converts the data into Oracle binary format in Hadoop, offloading the work to the cluster to minimize use of database CPU during the load.

Oracle Loader for Hadoop loads from a variety of formats. It can load data from text files, compressed files, sequence files, Hive tables, and data that can be accessed via a Hive table such as JSON, Parquet, ORC, and HBase.

Full documentation is available in the "Oracle Loader for Hadoop" chapter in the *Big Data Connectors User's Guide*.

You can initiate an Oracle Loader for Hadoop job by right-clicking on a database table icon in an Oracle Database connection. Note that the connection **must** be a service name connection when using direct path load. A SID connection will raise errors.

Using SQL Developer to run an Oracle Loader for Hadoop job requires an SSH connection to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster.

The file `.sqldev_olh_env` must exist in the home directory of the user of that SSH connection. The following is a sample `.sqldev_olh_env` file:

```
#!/bin/bash
export HIVE_HOME=/usr/lib/hive
export HIVE_CONF_DIR=${HIVE_HOME}/conf
export OLH_HOME=/u01/connectors/olh
export HADOOP_CLASSPATH=${OLH_HOME}/jlib/*:${HIVE_HOME}/lib/*:${HIVE_CONF_DIR}
export WALLET_LOCATION=/home/oracle/oracle_wallet
export TNS_ADMIN=${WALLET_LOCATION}
export OLH_LIB_JARS=${HIVE_HOME}/lib/hive-exec.jar,${HIVE_HOME}/lib/
libfb303-0.9.2.jar,${HIVE_HOME}/lib/hive-metastore.jar
export CLUSTER_WALLET_LOCATION=${WALLET_LOCATION}
export CLUSTER_TNS_ADMIN=${TNS_ADMIN}
```

1.18.3 Oracle SQL Connector for HDFS: Overview and Requirements

Oracle SQL Connector for HDFS can be used to:

- Load data into an Oracle Database table using SQL and external tables
- Query data in the Hadoop Distributed File System in-place via external tables

Oracle SQL Connector for HDFS supports

- Delimited text files
- Apache Hive tables over delimited text files
- Oracle Data Pump files generated by Oracle Loader for Hadoop

Oracle SQL Connector for HDFS creates an external table over delimited text files or a Hive table over delimited text files. This external table can be used to load or query in-place. (For more information, see the *Big Data Connectors User's Guide*.)

You can start Oracle SQL Connector for HDFS by right-clicking the Tables icon of an Oracle Database connection.

Using SQL Developer with Oracle SQL Connector for HDFS requires an SSH connection to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. The file `.sqldev_osch_env` must exist in the home directory of the user of that SSH connection. The following is a sample `.sqldev_osch_env` file:

```
#!/bin/bash
export HIVE_HOME=/usr/lib/hive
export HIVE_CONF_DIR=${HIVE_HOME}/conf
export OSCH_HOME=/u01/connectors/osch
export HADOOP_CLASSPATH=${OSCH_HOME}/jlib/*:${HIVE_HOME}/lib/*:${HIVE_CONF_DIR}
export WALLET_LOCATION=/home/oracle/oracle_wallet
export TNS_ADMIN=${WALLET_LOCATION}
export CLUSTER_WALLET_LOCATION=${WALLET_LOCATION}
export CLUSTER_TNS_ADMIN=${TNS_ADMIN}
```

1.19 SQL Developer Reports

SQL Developer provides many reports about the database and its objects. You can also create your own user-defined reports. To display reports, click the Reports tab on the left side of the window. If this tab is not visible, select **View** and then **Reports**. (As an alternative to the SQL Developer graphical interface, you can use the command-line interface for SQL Developer to generate reports.)

Individual reports are displayed in tabbed panes on the right side of the window; and for each report, you can select (in a drop-down control) the database connection for which to display the report. For reports about objects, the objects shown are only those visible to the database user associated with the selected database connection, and the rows are usually ordered by Owner. The detail display pane for a report includes the following icons at the top:

- **Freeze View** (the **pin**) keeps that report in the SQL Developer window when you click another report in the Reports navigator; a separate tab and detail view pane are created for that other report. If you click the pin again, the report's detail view pane is available for reuse.
- **Run Report** updates the detail view pane display by querying the database for the latest information.
- **Run Report in SQL Worksheet** displays the SQL statement used to retrieve the information for a report in a SQL Worksheet pane, where you can view, edit, and run the statement.

The time required to display specific reports will vary, and may be affected by the number and complexity of objects involved, and by the speed of the network connection to the database.

For most reports that contain names of database objects, you can double-click the object name in the report display pane (or right-click the object name and select Go To) to display that object in a detail view pane, just as if you had selected that object using the Connections navigator.

To export a report into an XML file that can be imported later, right-click the report name in the Reports navigator display and select **Export**. To import a report that had previously been exported, select the name of the report folder name (such as a user-defined folder) in which to store the imported report, right-click, and select **Import**.

You can create a **shared report** from an exported report by clicking Tools, then Preferences, and using the Database: User Defined Extensions pane to add a row with Type as REPORT and Location specifying the exported XML file. The next time you restart SQL Developer, the Reports navigator will have a Shared Reports folder containing that report.

Reports are grouped in the following categories:

[About Your Database reports \(page 1-84\)](#) list release information about the database associated with the connection.

[All Objects reports \(page 1-84\)](#) list information about all objects accessible to the user associated with the specified database connection, not just objects owned by the user.

[Application Express reports \(page 1-85\)](#) list information about Oracle Application Express 3.0.1 (or later) applications, pages, schemas, UI defaults, and workspaces.

[ASH and AWR reports \(page 1-85\)](#) list information provided by the Active Session History (ASH) and Automated Workload Repository (AWR) features.

[Database Administration reports \(page 1-85\)](#) list usage information about system resources.

[Data Dictionary reports \(page 1-85\)](#) list information about the data dictionary views that are accessible in the database. Examples of data dictionary views are ALL_OBJECTS and USER_TABLES.

[PL/SQL reports \(page 1-86\)](#) list information about your PL/SQL objects and allow you to search the source of those objects.

[Scheduler reports \(page 1-86\)](#) list information about jobs running on the database.

[Security reports \(page 1-86\)](#) list privilege-related information about the database.

[Streams reports \(page 1-87\)](#) list information about stream rules.

[Table reports \(page 1-87\)](#) list information about tables owned by the user associated with the specified connection. These reports can help you to better understand the metadata and data. The table reports include Quality Assurance reports that indicate possible logical design flaws and sources of runtime performance problems.

[XML reports \(page 1-89\)](#) list information about XML objects.

[Data Modeler reports \(page 1-89\)](#) list information about design objects that have been exported to the Data Modeler reporting repository.

TimesTen reports list information about Oracle TimesTen In-Memory database objects.

[User Defined reports \(page 1-89\)](#) are any customized reports that you have created.

1.19.1 Bind Variables for Reports

For some reports, you are prompted for **bind variables** before the report is generated. These bind variables enable you to further restrict the output. The default value for all bind variables is null, which implies no further restrictions. To specify a bind variable, select the variable name and type an entry in the Value field. Any bind variable values that you enter are case insensitive, all matches are returned where the value string appears anywhere in the name of the relevant object type.

1.19.2 About Your Database reports

The About Your Database reports list release information about the database associated with the selected connection. The reports include Version Banner (database settings) and National Language Support Parameters (NLS_xxx parameter values for globalization support).

1.19.3 All Objects reports

All Objects reports list information about objects visible to the user associated with the database connection.

All Objects: For each object, lists the owner, name, type (table, view, index, and so on), status (valid or invalid), the date it was created, and the date when the last data definition language (DDL) operation was performed on it. The Last DDL date can help you to find if any changes to the object definitions have been made on or after a specific time.

Collection Types: Lists information about for each collection type. The information includes the type owner, element type name and owner, and type-dependent specific information.

Dependencies: For each object with references to it, lists information about references to (uses of) that object.

Invalid Objects: Lists all objects that have a status of invalid.

Object Count by Type: For each type of object associated with a specific owner, lists the number of objects. This report might help you to identify users that have created an especially large number of objects, particularly objects of a specific type.

Public Database Links: Lists all public database links.

Public Synonyms: Lists all public synonyms.

1.19.4 Application Express reports

If you select a connection for a schema that owns any Oracle Application Express 3.0.1 (or later) applications, the Application Express reports list information about applications, pages, schemas, UI defaults, and workspaces. For information about Oracle Application Express, see the documentation for that product.

1.19.5 ASH and AWR reports

The ASH and AWR reports list information provided by the Active Session History (ASH) and Automated Workload Repository (AWR) features, which require special licensing.

1.19.6 Charts reports

Charts reports include a chart showing the distribution of objects of various object types (number of tables, indexes, and so on).

1.19.7 Database Administration reports

Database Administration reports list usage information about system resources. This information can help you to manage storage, user accounts, and sessions efficiently. (The user for the database connection must have the DBA role to see most Database Administration reports.)

All Tables: Contains the reports that are also grouped under Table reports, including Quality Assurance reports.

Cursors: Provide information about cursors, including cursors by session (including open cursors and cursor details).

Database Parameters: Provide information about all database parameters or only those parameters that are not set to their default values.

Locks: Provide information about locks, including the user associated with each.

Sessions: Provide information about sessions, selected and ordered by various criteria.

Storage: Provide usage and allocation information for tablespaces and data files.

Top SQL: Provide information about SQL statements, selected and ordered by various criteria. This information might help you to identify SQL statements that are being executed more often than expected or that are taking more time than expected.

Users: Provide information about database users, selected and ordered by various criteria. For example, you can find out which users were created most recently, which user accounts have expired, and which users use object types and how many objects each owns.

Waits and Events: Provide information about waits and events, selected by criteria related to time and other factors. For Events in the Last x Minutes, specify the number of minutes in the Enter Bind Values dialog box.

1.19.8 Data Dictionary reports

Data Dictionary reports list information about the data dictionary views that are accessible in the database. Examples of data dictionary views are ALL_OBJECTS and USER_TABLES.

Dictionary View Columns: For each Oracle data dictionary view, lists information about the columns in the view.

Dictionary Views: Lists each Oracle data dictionary view and (in most cases) a comment describing its contents or purpose.

1.19.9 Scheduler reports

Jobs reports list information about jobs and other objects related to scheduling jobs using SQL Developer.

DBMS Jobs: Lists information about all jobs, DBA jobs (jobs for which a DBA user is associated with the database connection), and your jobs (jobs for which the user associated with the database connection is the log user, privilege user, or schema user). The information for each job includes the start time of its last run, current run, and next scheduled run.

Definitions: Lists information about all objects of types associated with job scheduling.

Executions: Lists information about the executions of jobs.

1.19.10 PL/SQL reports

PL/SQL reports list information about PL/SQL packages, function, and procedures, and about types defined in them.

Program Unit Arguments: For each argument (parameter) in a program unit, lists the program unit name, the argument position (1, 2, 3, and so on), the argument name, and whether the argument is input-only (In), output-only (Out), or both input and output (In/Out).

Search Source Code: For each PL/SQL object, lists the source code for each line, and allows the source to be searched for occurrences of the specified variable.

Unit Line Counts: For each PL/SQL object, lists the number of source code lines. This information can help you to identify complex objects (for example, to identify code that may need to be simplified or divided into several objects).

1.19.11 Security reports

Security reports list information about users that have been granted privileges, and in some cases about the users that granted the privileges. This information can help you (or the database administrator if you are not a DBA) to understand possible security issues and vulnerabilities, and to decide on the appropriate action to take (for example, revoking certain privileges from users that do not need those privileges).

Auditing: Lists information about audit policies.

Encryption: Lists information about encrypted columns.

Grants and Privileges: Includes the following reports:

- **Column Privileges:** For each privilege granted on a specific column in a specific table, lists the user that granted the privilege, the user to which the privilege was granted, the table, the privilege, and whether the user to which the privilege was granted can grant that privilege to other users.
- **Object Grants:** For each privilege granted on a specific table, lists the user that granted the privilege, the user to which the privilege was granted, the table, the

privilege, and whether the user to which the privilege was granted can grant that privilege to other users.

- **Role Privileges:** For each granted role, lists the user to which the role was granted, the role, whether the role was granted with the ADMIN option, and whether the role is designated as a default role for the user.
- **System Privileges:** For each privilege granted to the user associated with the database connection, lists the privilege and whether it was granted with the ADMIN option.

Policies: Lists information about policies.

Public Grants: Lists information about privileges granted to the PUBLIC role.

1.19.12 Streams reports

Streams reports list information about stream rules.

All Stream Rules: Lists information about all stream rules. The information includes stream type and name, rule set owner and name, rule owner and name, rule set type, streams rule type, and subsetting operation.

Your Stream Rules: Lists information about each stream rule for which the user associated with the database connection is the rule owner or rule set owner. The information includes stream type and name, rule set owner and name, rule owner and name, rule set type, streams rule type, and subsetting operation.

1.19.13 Table reports

Table reports list information about tables owned by the user associated with the specified connection. This information is not specifically designed to identify problem areas; however, depending on your resources and requirements, some of the information might indicate things that you should monitor or address.

For table reports, the owner is the user associated with the database connection.

Columns: For each table, lists each column, its data type, and whether it can contain a null value. Also includes **Datatype Occurrences:** For each table owner, lists each data type and how many times it is used.

Comments for tables and columns: For each table and for each column in each table, lists the descriptive comments (if any) associated with it. Also includes a report of tables without comments. If database developers use the COMMENT statement when creating or modifying tables, this report can provide useful information about the purposes of tables and columns

Constraints: Includes the following reports related to constraints:

- **All Constraints:** For each table, lists each associated constraint, including its type (unique constraint, check constraint, primary key, foreign key) and status (enabled or disabled).
- **Check Constraints:** For each check constraint, lists information that includes the owner, the table name, the constraint name, the constraint status (enabled or disabled), and the constraint specification.
- **Enabled Constraints and Disabled Constraints:** For each constraint with a status of enabled or disabled, lists the table name, constraint name, constraint type (unique constraint, check constraint, primary key, foreign key), and status. A

disabled constraint is not enforced when rows are added or modified; to have a disabled constraint enforced, you must edit the table and set the status of the constraint to Enabled.

- **Foreign Key Constraints:** For each foreign key constraint, lists information that includes the owner, the table name, the constraint name, the column that the constraint is against, the table that the constraint references, and the constraint in the table that is referenced.
- **Primary Key Constraints:** For primary key constraint, lists information that includes the owner, the table name, the constraint name, the constraint status (enabled or disabled), and the column name.
- **Unique Constraints:** For each unique constraint, lists information that includes the owner, the table name, the constraint name, the constraint status (enabled or disabled), and the column name.

Indexes: Includes information about all indexes, indexes by status, indexes by type, and unused indexes.

Organization: Specialized reports list information about partitioned tables, clustered tables, and index-organized tables.

Quality Assurance: Reports that identify conditions that are not technically errors, but that usually indicate flaws in the database design.

Statistics: For each table, lists statistical information, including when it was last analyzed, the total number of rows, the average row length, and the table type. In addition, specialized reports order the results by most rows and largest average row length.

Storage: Lists information about the table count by tablespace and the tables in each tablespace.

Triggers: Lists information about all triggers, disabled triggers, and enabled triggers.

User Synonyms: Displays information about either all user synonyms or those user synonyms containing the string that you specify in the Enter Bind Variables dialog box (uncheck Null in that box to enter a string).

User Tables: Displays information about either all tables or those tables containing the string that you specify in the Enter Bind Variables dialog box (uncheck Null in that box to enter a string).

1.19.13.1 Quality Assurance reports

Quality assurance reports are table reports that identify conditions that are not technically errors, but that usually indicate flaws in the database design. These flaws can result in various problems, such as logic errors and the need for additional application coding to work around the errors, as well as poor performance with queries at runtime.

Tables without Primary Keys: Lists tables that do not have a primary key defined. A primary key is a column (or set of columns) that uniquely identifies each row in the table. Although tables are not required to have a primary key, it is strongly recommended that you create or designate a primary key for each table. Primary key columns are indexed, which enhances performance with queries, and they are required to be unique and not null, providing some "automatic" validation of input data. Primary keys can also be used with foreign keys to provide referential integrity.

Tables without Indexes: Lists tables that do not have any indexes. If a column in a table has an index defined on it, queries that use the column are usually much faster and more efficient than if there is no index on the column, especially if there are many rows in the table and many different data values in the column.

Tables with Unindexed Foreign Keys: Lists any foreign keys that do not have an associated index. A foreign key is a column (or set of columns) that references a primary key: that is, each value in the foreign key must match a value in its associated primary key. Foreign key columns are often joined in queries, and an index usually improves performance significantly for queries that use a column. If an unindexed foreign key is used in queries, you may be able to improve runtime performance by creating an index on that foreign key.

1.19.14 XML reports

XML reports list information about XML objects.

XML Schemas: For each user that owns any XML objects, lists information about each object, including the schema URL of the XSD file containing the schema definition.

1.19.15 Data Modeler reports

Data Modeler reports list information about objects that have been exported to the Data Modeler **reporting repository**. (To export to the reporting repository, you must use the full Data Modeler product; you cannot use the Data Modeler that is included in SQL Developer.)

To use any Data Modeler reports in SQL Developer, you must have created a database connection to the schema that contains the Data Modeler reporting repository.

To view any Data Modeler report, expand the Data Modeler Reports hierarchy and double-click the desired report. If you are prompted for a database connection, select the one containing the reporting repository. If you are prompted for bind variables, select the variable name and type an entry in the Value field. Any bind variable values that you enter are case insensitive, all matches are returned where the value string appears anywhere in the name of the relevant object type.

Design Content: Includes reports about the data types and the logical and relational models.

Design Rules: Includes reports about potential violations of the Data Modeler design rules. For example, for logical model attributes, you can see attributes without data types and attributes not based on domains; and for relational model tables, you can see tables without columns, primary keys, foreign key links, and unique constraints.

See Also:

For information about exporting designs to the reporting repository, see *Oracle SQL Developer Data Modeler User's Guide*.

1.19.16 User Defined reports

User Defined reports are any reports that are created by SQL Developer users. To create a user-defined report, right-click the User Defined node under Reports and select **New Report**. A dialog box is displayed in which you specify the report name and the SQL query to retrieve information for the report.

You can organize user-defined reports in folders, and you can create a hierarchy of folders and subfolders. To create a folder for user-defined reports, right-click the User Defined node or any folder name under that node and select **Add Folder**.

Information about user-defined reports, including any folders for these reports, is stored in a file named UserReports.xml under the directory for user-specific information.

For examples of creating user-defined reports, see:

- [User-Defined Report Example: Chart](#) (page 1-90)
- [User-Defined Report Example: Dynamic HTML](#) (page 1-91)
- [Gauges: In the SQL Worksheet and User-Defined Reports](#) (page 1-49)

1.19.16.1 User-Defined Report Example: Chart

This example creates a report displayed as a chart. It uses the definition of the EMPLOYEES table from the HR schema, which is a supplied sample schema.

Right-click on User Defined Reports and select **New Report**. In the Create Report dialog box, specify a report name; and for **SQL**, enter the following:

```
select m.department_id, e.last_name, e.salary
from employees m, employees e
where e.employee_id = m.employee_id
order by 1
```

The preceding query lists the last name and salary of each employee in each department, grouping the results by department ID (10, 20, 30, ... 110). Note that the expected syntax for the SQL statement for a chart report is:

```
SELECT <group>,<series>,<value> FROM <table(s)>
```


Click the **Property** subtab and specify the following:

- **Chart Type: Bar - Vertical Stack**
- (Default values for remaining options)

Click **Apply**.

Use the Reports navigator to view the newly created user-defined report. For **Connection**, specify one that connects to the HR sample schema.

The report is displayed as a chart, part of which is shown in the following illustration. For example, as you can see, department 50 has mainly employees with the lowest salaries, and department 90 consists of the three highest-paid employees.



1.19.16.2 User-Defined Report Example: Dynamic HTML

This example creates a report using one or more PL/SQL DBMS_OUTPUT statements, so that the report is displayed as dynamic HTML.

Right-click on User Defined Reports and select **New Report**. In the New Report dialog box, specify a report name; for **Style**, select **plssql-dbms_output**; and for **SQL**, enter the following:

```
begin
dbms_output.put_line ('<H1> This is Level-1 Heading </H1>');
dbms_output.put_line ('<H2> This is a Level-2 Heading </H2>');
dbms_output.put_line ('<p> This is regular paragraph text. </p>');
end;
```

Click **Apply**.

Use the Reports navigator to view the newly created user-defined report. For **Connection**, specify any from the list. (This report does not depend on a specific connection or table.).

The report is displayed as formatted HTML output.

1.20 SQL Developer Preferences

You can customize many aspects of the SQL Developer interface and environment by modifying SQL Developer preferences according to your preferences and needs. To modify SQL Developer preferences:

- Windows and Linux systems: Select **Tools**, then **Preferences**.
- Mac OS X systems: Use the **Oracle SQL Developer** menu to access the Preferences dialog.

Information about SQL Developer preferences is stored under the directory for user-specific information.

Most preferences are self-explanatory, and this topic explains only those whose meaning and implications are not obvious. Some preferences involve performance or system resource trade-offs (for example, enabling a feature that adds execution time), and other preferences involve only personal aesthetic taste. The preferences are grouped in the following categories.

Search box: You can enter a string to limit the tree display to matching relevant preference groups.

1.20.1 Environment

The Environment pane contains options that affect the startup and overall behavior and appearance of SQL Developer. You can specify that certain operations be performed automatically at specified times, with the trade-off usually being the extra time for the operation as opposed to the possibility of problems if the operation is not performed automatically (for example, if you forget to perform it when you should).

The undo level (number of previous operations that can be undone) and navigation level (number of open files) values involve slight increases or decreases system resource usage for higher or lower values.

Save All When Deactivating or Exiting: If this option is checked, any changes to your files are automatically saved when you close or stop SQL Developer. If this option is not checked, you are asked if you want to save changes (if there are any).

Automatically Reload Externally Modified Files: If this option is checked, any files open in SQL Developer that have been modified by an external application are updated when you switch back to SQL Developer, overwriting any changes that you might have made. If this option is not checked, changes that you make in SQL Developer overwrite any changes that might have been made by external applications.

Silently Reload When File Is Unmodified: If this option is checked, you are not asked if you want to reload files that have been modified externally but not in SQL Developer. If this option is not checked, you are asked if you want to reload each file that has been modified externally, regardless of whether it has been modified in SQL Developer.

Undo Level: The number of undo operations that SQL Developer will remember.

Navigation Level: The the number of times you can go back and forth on browsing history by using the Back and Forward navigation buttons.

Look and Feel: A JGoodies look and feel. The new look and feel will be reflected after you restart SQL Developer.

Line Terminator: Style for line termination. This line termination style will be applied to all new files.

Encoding: Character set to be used for encoding. All new files will be created with the new encoding, and the new encoding will be used to open and save files that do not have encoding tags.

Reset Skipped Messages: Causes all message types to be displayed in the future, including those that you have previously opted not to see.

Environment: Local History

The Local History pane controls whether information about editing operations on files opened within SQL Developer is kept. If local history is enabled, you can specify how long information is retained and the maximum number of revisions for each file.

Environment: Log

The Log pane configures the colors of certain types of log messages and the saving of log messages to log files.

Save Logs to File: If this option is checked, all output to the Messages - Log window is saved to log files, where the file name reflects the operation and a timestamp. You are also asked to specify a **Log Directory**; and if the specified directory does not already exist, it is created if and when it is needed. Note that if you save log information to files, the number of these files can become large.

Maximum Log Lines: The maximum number of lines to store in each log file.

1.20.2 Change Management Parameters

The Change Management Parameters pane contains options that affect Change Manager support in SQL Developer, especially how objects are compared to determine if a change has occurred and how extensive the DDL statements are when the script for deploying changes is generated.

Ignore physical attributes: Causes differences in physical attributes to be ignored when objects are compared.

Ignore tablespace: Causes differences in the tablespace specified for the object to be ignored when objects are compared.

Ignore segment attributes: Causes differences in segment specification to be ignored when objects are compared.

Ignore storage: Causes differences in storage specification to be ignored when objects are compared.

Ignore table column positions: Causes differences in the positions of table columns to be ignored when objects are compared. For example, if this option is enabled (checked), two tables would not be considered different in the only difference is that Column1 and Column2 appear first and second in one table but second and first in the other table.

Match Constraints: Controls whether constraints are matched for comparison by the names or definitions or the constraints (**Match constraints by name** or **Match constraints by definition**).

Report constraint name difference: If Match constraints by definition is selected, causes the constraint name to be displayed when differences are reported.

Generate constraint indexes: Causes constraint indexes to be generated in the DDL for the deployment script.

Generate constraints: Causes constraints to be generated in the DDL for the deployment script.

Generate password values: Causes password values to be generated in the DDL for the deployment script.

Generate referential constraints: Causes referential constraints to be generated in the DDL for the deployment script.

Generate SQL terminator: Causes the SQL statement terminator character to be generated at the end of each statement in the DDL for the deployment script.

Change Management Parameters: Conflict Resolution

The Conflict Resolution pane contains options for resolving conflicts when the old or existing item's value is different from the potential new or modified value or when the existing item would be removed.

Item to be added already exists, or the item to be removed is different. Apply Exists Different: FALSE causes the item to be added or the existing item to be removed; TRUE causes the item not to be added or the existing item not to be removed.

Item to be modified has already been modified to a different value. Apply Modified Different: FALSE causes the item to be modified so as to replace the existing (different) value; TRUE causes the item to keep the existing (different) value.

Item to be modified has already been modified to the old value recorded in the difference. Apply Modified Previously: FALSE causes the item to be changed to the new value; TRUE causes the item to keep the existing value (the old value recorded in the difference).

1.20.3 Code Editor

The Code Editor pane contains general options that affect the appearance and behavior of SQL Developer when you edit functions, procedures, and packages.

Autopin PL/SQL Editors: Keeps the current PL/SQL editor open when you open another function, procedure, or package.

Max Open PL/SQL Editors: Specifies the maximum number of PL/SQL editors that can be kept open ("pinned").

Start in Read Only Mode: Causes functions, procedures, and packages to be displayed initially in read-only mode. (To switch to read/write mode, click the Switch to Write Mode icon in the code editing window.)

Link Stored Procedures to Files: If this option is enabled (checked), then every time you compile a file-based PL/SQL module, SQL Developer associates the database source with the file-based one. Subsequently, during a debugging session, whenever a reference to a PL/SQL module is made, SQL Developer checks its internal database and file-based source associations; and if it is able to find a file-based source with the specified name, it opens it. This setting is of interest to users who keep their PL/SQL code in a source control system.

If this option is disabled, then during a debugging session SQL Developer always opens the procedure residing in the database, and does not check the file system.

Auto-Indent New Lines: Automatically indents a new line when you press Enter at the end of a line. The new line will automatically be indented at the same initial indentation as the line preceding it.

Perform Block Indent or Outdent for Selections: Performs a block indent or block outdent on a selection when you press Tab or Shift+Tab, respectively. With this option selected, when you press Tab on a selected block of text, the entire block will be indented to the current tab size. Shift+Tab on the same block would outdent it, as a block, to the current tab size.

Use Smart Home: Contextualizes the cursor's understanding of home (the beginning of the line). With this setting selected, pressing Home positions the cursor at the start of the line after any leading spaces or tabs. Pressing Home again repositions the cursor at the start of the line before any leading spaces or tabs. Continuing to press Home toggles the cursor between these two locations.

With this setting deselected, pressing Home simply places the cursor at the start of the line.

Use Smart End: Contextualizes the cursor's understanding of end of line. The behavior is analogous to that for Smart Home, except that the cursor responds to the End key, and its behavior regarding the end of the line and any trailing spaces is altered.

Use Jump Scrolling for Keyboard Navigation: Implement jump scrolling, which involves behavior of the keyboard arrow keys. With this setting selected, when you navigate off-screen using the keyboard arrow keys, the editor view will "jump" to recenter the cursor location in the middle of the editor view.

With this setting deselected, the editor view will scroll the editor view the minimum amount to bring the cursor back into view.

Use Change of Case As Word Boundary: Has change of case regarded as the boundary of a word, for example, when you double-click to select a word.

Enable Cut or Copy of Current Line with No Selection: Applies all cut and copy operations to the current line whenever there is no text selection in the editor.

Automatically Copy Paste Imports: Automatically add imports when references are introduced to objects that have not yet been imported.

Adjust Indentation When Pasting: Corrects the indentation of a pasted in item that includes indentation.

Escape When Pasting in String Literals: Includes the correct escape characters in pasted-in string literals.

Supported Gutter Navigation: Enables the use of blue arrows to go from procedure definition to declaration and back.

Supported Quick Documentation: Enables the use of Ctrl+D to show the Javadoc associated with the object or referenced object (if there is any associated Java method) at the current cursor location. (Although Javadoc-style comments are not used in PL/SQL, the signature of the method and any associated comments may be useful.)

Supported Sync Spec and Body: Enables you to right-click in a package body and select Synchronize Specification and Body, which updates body to reflect any changes to the specification (for example, to add the implement new methods for declarations added to the specification or to add declarations to the specification for private methods in the body to then make them public).

Supported Test Query: Enables you to right-click and select Test Query when the cursor is within a SELECT statement, to check to that the correct rows are returned.

Code Editor: Bookmarks

The Bookmarks pane contains options that determine the persistence and search behavior for bookmarks that you create when using the code editor.

Code Editor: Caret Behavior

The Caret Behavior pane contains options that determine the shape, color, and blinking characteristics of the caret (cursor) in the code editor.

Code Editor: Completion Insight

The Completion Insight pane contains options for the logical completion (autocomplete options) of keywords and names while you are coding in the SQL Worksheet.

When you pause for the auto-popup time (if the auto-popup is enabled) or when you press **Ctrl+Space**, code insight provides a context-sensitive popup window that can help you select parameter names. Completion insight provides you with a list of possible completions at the insertion point that you can use to auto-complete code you are editing. This list is based on the code context at the insertion point. To exit code insight at any time, press Esc or continue typing.

You can enable or disable automatic completion and parameter insight, as well as set the time delay for the popup windows.

Suppress automatic pop-ups for more than [number] matches: If you enable completion auto-popup in either or both of the SQL Worksheet and PL/SQL Editor, you can limit popups to situations where the number of available suggestions is less than or equal to the specified value. (Smaller values will generally result in fewer popups while you are typing.) This option does not apply when you explicitly invoke completion by pressing Ctrl+Space.

Generate Column/Table Aliases Automatically: Automatically generates table aliases if you select multiple tables from the popup window; and if you then edit the column list, each column name in the popup window is prefixed with a table alias.

Change case as you type: If you enable this option, select the desired action: **None** makes no changes to what you type (same effect as not enabling the option); **Initial Cap** capitalizes the first letter of each word; **Upper Case** capitalizes all letters; **Upper Keywords** capitalizes all keywords; **Lower Case** makes each letter lowercase; **Lower Keywords, Upper Identifiers** lowercases all keywords but capitalizes all identifiers.

Enable Semantic Analysis Info Tip: Lets you get advice on "grammatical" issues with SQL statements. For a detailed explanation, see <http://www.thatjeffsmith.com/archive/2013/07/semantic-analysis-info-tips-in-oracle-sql-developer-4-0/>.

Autogenerate GROUP BY Clause: Automatically generates a GROUP BY clause if you manually enter (not copy/paste) a SELECT statement containing a COUNT function, and then edit the SELECT query.

Code Editor: Display

The Display pane contains general options for the appearance and behavior of the code editor.

Enable Text Anti-Aliasing allows smooth-edged characters where possible.

Show Whitespace Characters renders spaces, new lines, carriage returns, non-breaking spaces, and tab characters as alternate visible characters.

Show Breadcrumbs shows the breadcrumb bar, which shows the hierarchy of nodes from the current caret position up to the top of the file. Hover the mouse cursor over a node to display information about the node.

Show Scroll Tip enables the Tip window that displays the methods in view while scrolling.

Show Code Folding Margin allows program blocks in procedures and functions to be expanded and collapsed in the display.

Show visible Right Margin renders a right margin that you can set to control the length of lines of code.

Enable Automatic Brace Matching controls the highlighting of opening parentheses and brackets and of blocks when a closing parenthesis or bracket is typed.

Code Editor: Fonts

The Fonts pane specifies text font options for the code editor.

Display Only Fixed-Width Fonts: If this option is checked, the display of available font names is restricted to fonts where all characters have the same width. (Fixed-width fonts are contrasted with proportional-width fonts.)

Code Editor: Format

The Format pane controls how statements in the SQL Worksheet are formatted when you click Format SQL. The options include whether to insert space characters or tab characters when you press the Tab key (and how many characters), uppercase or lowercase for keywords and identifiers, whether to preserve or eliminate empty lines, and whether comparable items should be placed on the same line (if there is room) or on separate lines.

The Advanced Format subpane lets you specify more detailed formatting options. It also includes these options:

- **Preview with Current Settings:** You specify changes on the left side, and the preview area on the right side reflects the changes.
- **Auto-Detect Formatter Settings:** You paste code with the desired formatting into the preview pane on the right side, and SQL Developer adjusts the settings on the left side to reflect what you pasted. (It automatically detects, or autodetects, your setting preferences.)

You can export these settings to a code style profile XML file, and can import settings from a previously exported code style profile file.

The Custom Format subpane lets you further customize your formatting settings. You can also export these settings to a custom formatter program XML file, and you can import settings from a previously exported custom formatter program file.

Code Editor: Line Gutter

The Line Gutter pane specifies options for the line gutter (left margin of the code editor).

Show Line Numbers: If this option is checked, lines are numbered. (To go to a line number while you are using the SQL Worksheet, press Ctrl+G.)

Enable Line Selection by Click-Dragging: If this option is checked, you can select consecutive lines in the editor by clicking in the gutter and dragging the cursor without releasing the mouse button.

Code Editor: PL/SQL Syntax Colors

The PL/SQL Syntax Colors pane specifies colors for different kinds of syntax elements.

Code Editor: Printing

The Printing pane specifies options for printing the contents of the code editor. The Preview pane sample display changes as you select and deselect options.

Code Editor: Printing HTML

The Printing HTML pane specifies options for printing HTML files from the code editor.

Code Editor: Undo Behavior

The Undo Behavior pane specifies options for the behavior of undo operations (Ctrl +Z, or Edit, then Undo). Only consecutive edits *of the same type* are considered; for example, inserting characters and deleting characters are two different types of operation.

Allow Navigation-Only Changes to be Undoable: If this option is checked, navigation actions with the keyboard or mouse can be undone. If this option is not checked, navigation actions cannot be undone, and only actual changes to the text can be undone.

1.20.4 Compare and Merge

The Compare and Merge panes define options for comparing and merging two source files.

- Compare Preferences
- Merge Preferences
- Comparing Source Files

Compare Preferences

Ignore Whitespace: If this option is enabled, leading and trailing tabs and letter spacing are ignored when comparing files. Carriage returns are not ignored. Enabling this option makes comparing two files easier when you have replaced all the space with hard tabs, or vice versa. Otherwise, every line in the two documents might be shown as different in the Compare window.

Show Character Differences: If this option is enabled, characters that are present in one file and not in another are highlighted. Red highlighting indicates a character that has been removed. Green highlighting indicates a character that has been added. The highlighting is shown only when you click into a comparison block that contains character differences.

Maximum File Size (KB): the maximum file size (number of kilobytes) for which the operation will be performed. You can specify different sizes for text files and XML files.

Enable Java Compare: If this option is enabled, Java source files can be compared in a structured format.

Enable XML Compare: If this option is enabled, XML files can be compared.

Merge Preferences

Enable XML Merge: If this option is enabled, XML files can be merged.

Reformat Result: If this option is enabled, merged XML files can be reformatted.

Validate Result (May require Internet access): If this option is enabled, merged XML files will be validated.

Maximum File Size (KB): the maximum XML file size (number of kilobytes) for which the operation will be performed.

Comparing Source Files

You can compare source files in the following ways:

- A file currently being edited with its saved version: Place the focus on the current version open in the editor; from the main menu, select **File**, then **Compare With**, then **File on Disk**.
- One file with another file: Place the focus on the file in the editor to be compared; from the main menu, select **File**, then **Compare With**, then **Other File**. In the Select File to Compare With dialog, navigate to the file and click Open.

1.20.5 Database

The Database pane sets properties for the database connection.

Validate date and time default values: If this option is checked, date and time validation is used when you open tables.

Filename for connection startup script: File name for the startup script to run when an Oracle database connection is opened. You can click **Browse** to specify the location. The default location is the default path for scripts (see the Database: Worksheet preferences pane).

Database: Advanced

The Advanced pane specifies options such as the **SQL Array Fetch Size** and **Display** options for null values and for STRUCT object values.

You can also specify Kerberos thin driver configuration parameters, which enables you to create database connections using Kerberos authentication and specifying the user name and password. For more information, see the Kerberos Authentication explanation on the Oracle tab in the [Create/Edit/Select Database Connection](#) (page 6-11) dialog box. For information about configuring Kerberos authentication, see *Oracle Database Security Guide*.

SQL Array Fetch Size: Specifies the maximum number of rows to be fetched at a time into the table Data grid display.

Display Null Value As: Specifies the value to be displayed in the table Data grid for null values. The default is the string (null).

Display Null Using Background Color: Specifies the background color for any null values displayed in the table Data grid.

Display Struct Value in Grid: Controls how unstructured data types are displayed in the table Data grid. For example, if this option is checked, Oracle Spatial and Graph geometry data is displayed with the full SDO_GEOMETRY definition; but if this option is not checked, each Oracle Spatial and Graph geometry object is displayed as just [MDSYS.SDO_GEOMETRY].

Autocommit: If this option is checked, a commit operation is automatically performed after each INSERT, UPDATE, or DELETE statement executed using the SQL Worksheet. If this option is not checked, a commit operation is not performed until you execute a COMMIT statement.

Kerberos Thin Config: Config File: Kerberos configuration file (for example, krb5.conf). If this is not specified, default locations will be tried for your Java and system configuration.

Kerberos Thin Config: Credential Cache File: Kerberos credential cache file (for example, krb5_cc_cache). If this is not specified, a cache will not be used, and a principal name and password will be required each time.

Use Oracle Client: Lets you specify the use of an Oracle Home or Oracle Instant Client. Click **Configure** to select the **Client Type** (Oracle Home or Instant Client) and the **Client Location**.

Use OCI/Thick driver: If this option is checked, and if an OCI (thick, Type 2) driver is available, that driver will be used instead of a JDBC (thin) driver for basic and TNS (network alias) database connections. If any connections use a supported Remote Authentication Dial In User Service (RADIUS) server, check this option.

Tnsnames Directory: Enter or browse to select the location of the tnsnames.ora file. If no location is specified, SQL Developer looks for this file as explained in [Database Connections](#) (page 1-28). Thus, any value you specify here overrides any TNS_ADMIN environment variable or registry value or (on Linux systems) the global configuration directory.

Database: Autotrace/Explain Plan

The Autotrace/Explain Plan pane specifies information to be displayed on the Autotrace and Explain Plan panes in the SQL Worksheet.

Database: Drag and Drop

The Drag and Drop Effects pane determines the type of SQL statement created in the SQL Worksheet when you drag an object from the Connections navigator into the SQL Worksheet. The SQL Developer preference sets the default, which you can override in the Drag and Drop Effects dialog box.

The type of statement (INSERT, DELETE, UPDATE, or SELECT) applies only for object types for which such a statement is possible. For example, SELECT makes sense for a table, but not for a trigger. For objects for which the statement type does not apply, the object name is inserted in the SQL Worksheet.

Database: Instance Viewer

The Instance Viewer pane determines default behaviors when a DBA user uses the database Instance Viewer.

Graph Duration: The amount of history shown in graphs that include a time axis. For example, you might want graphs to reflect data only for the last two minutes or for the last hour.

Database: Licensing

Some SQL Developer features require that licenses for specific Oracle Database options be in effect for the database connection that will use the feature. The Licensing pane enables you to specify, for each defined connection, whether the database has the Oracle Change Management Pack, the Oracle Tuning Pack, and the Oracle Diagnostics Pack.

For each cell in this display (combination of license and connection), the value can be true (checked box), false (cleared box), or unspecified (solid-filled box).

If an option is specified as true for a connection in this pane, you will not be prompted with a message about the option being required when you use that connection for a feature that requires the option.

Database: Navigation Filter

The Navigation Filter pane enables you to control the types of objects that appear in the Connections navigator display for connections to Oracle and third-party databases.

If you check **Enable Navigation Tree Filtering**, you can click separate tabs for different database types, and for each type you can specify (by check marks) which types of objects appear in the navigator hierarchy for connections to that type of database.

Database: NLS

The NLS pane specifies values for globalization support parameters, such as the language, territory, sort preference, and date format. These parameter values are used for SQL Developer session operations, such as for statements executed using the SQL Worksheet and for the National Language Support Parameters report. Specifying values in this preferences pane does *not* apply those values to the underlying database itself. To change the database settings, you must change the appropriate initialization parameters and restart the database.

Note that SQL Developer does not use default values from the current system for globalization support parameters; instead, SQL Developer, when initially installed, by default uses parameter values that include the following:

```
NLS_LANG, "AMERICAN"
NLS_TERR, "AMERICA"
NLS_CHAR, "AL32UTF8"
NLS_SORT, "BINARY"
NLS_CAL, "GREGORIAN"
NLS_DATE_LANG, "AMERICAN"
NLS_DATE_FORM, "DD-MON-RR"
```

Database: ObjectViewer Parameters

The ObjectViewer Parameters pane specifies whether to freeze object viewer windows, whether to open objects with a single click instead of a double click, and display options for the output. The display options will affect the generated DDL on the SQL tab. The Data Editor Options affect the behavior when you are using the Data tab to edit table data.

Data Editor Options

Post Edits on Row Change: If this option is checked, posts DML changes when you perform edits using the Data tab (and the Set Auto Commit On option determines whether or not the changes are automatically committed). If this option is not checked, changes are posted and committed when you press the Commit toolbar button.

Set Auto Commit On (available only if Post Edit on Row Changes is enabled): If this option is checked, DML changes are automatically posted and committed when you perform edits using the Data tab.

Clear persisted table column widths, order, sort, and filter settings: If you click **Clear**, then any customizations in the Data tab display for table column widths, order, sort, and filtering are not saved for subsequent openings of the tab, but instead the default settings are used for subsequent openings.

Use ORA_ROWSCN for DataEditor insert and update statements: If this option is checked, SQL Developer internally uses the ORA_ROWSCN pseudocolumn in performing insert and update operations when you use the Data tab. If you experience any errors trying to update data, try *unchecking* (disabling) this option.

Database: PL/SQL Compiler

The PL/SQL Compiler pane specifies options for compilation of PL/SQL subprograms.

Types of messages: You can control the display of informational, severe, and performance-related messages. (The ALL type overrides any individual specifications for the other types of messages.) For each type of message, you can specify any of the following:

- **No entry (blank):** Use any value specified for ALL; and if none is specified, use the Oracle default.
- **Enable:** Enable the display of all messages of this category.
- **Disable:** Disable the display of all messages of this category.
- **Error:** Enable the display of only error messages of this category.

Optimization Level: 0, 1, or 2, reflecting the optimization level that will be used to compile PL/SQL library units. The higher the setting of this parameter, the more effort the compiler makes to optimize PL/SQL library units. However, for a module to be compiled with PL/SQL debugging information, the level must be 0 or 1.

PLScope Identifiers: Specifies the amount of PL/Scope identifier data to collect and use (All or None).

Drop Type Force: If this option is checked, the FORCE option is used if you attempt to delete an object type that has dependencies, and the type is deleted. If this option is not checked, object types with dependencies cannot be deleted. (You should generally not delete object types that have dependencies.)

Database: Reports

The Reports pane specifies options relating to SQL Developer reports.

Close all reports on disconnect: If this option is checked, all reports for any database connection are automatically closed when that connection is disconnected.

Chart Row Limit: Lets you specify a maximum number of rows to display for chart reports.

Database: SQL Editor Code Templates

The SQL Editor Code Templates pane enables you to view, add, and remove templates for editing SQL and PL/SQL code. Code templates assist you in writing code more quickly and efficiently by inserting text for commonly used statements. You can then modify the inserted text.

The template ID string is not used by SQL Developer; only the template content (Description text) is used, in that it is considered by completion insight (explained in Code Editor: **Completion Insight**) in determining whether a completion popup should be displayed and what the popup should contain. For example, if you define code template ID *mydate* as *SELECT sysdate FROM dual*, then if you start typing *select* in the SQL Worksheet, the auto-popup includes *SELECT sysdate FROM dual*.

Add Template: Adds an empty row in the code template display. Enter an ID value, then move to the Template cell; you can enter template content in that cell, or click the pencil icon to open an editing box to enter the template content.

Remove Template: Deletes the selected code template.

Import: Lets you import code style profile settings that you previously exported.

Export: Exports the current code profile settings to an XML file.

Autoformat PL/SQL in Procedures, Packages, Views, and Triggers: If this option is checked, the Format options are applied automatically as you enter and modify PL/SQL code in procedures, packages, views, and triggers; if this option is not checked, the Format options are applied only when you so request.

Database: Third Party JDBC Drivers

The Third Party JDBC Drivers pane specifies drivers to be used for connections to MySQL and third-party (non-Oracle) databases, such as Amazon Redshift, IBM DB2, Microsoft SQL Server, Sybase Adaptive Server, PostgreSQL, or Apache Hive. To add a driver, click **Add Entry** and select the path for the driver:

- For Amazon Redshift: Download from <http://docs.aws.amazon.com/redshift/latest/mgmt/configure-jdbc-connection.html#download-jdbc-driver>. The file name is RedshiftJDBC4n-1.2.1.1001.jar, where *n* is 0, 1, or 2 depending on the version of the JDBC API that SQL Developer uses.
- For Hive, see the information and downloads available at <http://www.cloudera.com>. Using Cloudera_HiveJDBC4_*.zip is recommended, for example, Cloudera_HiveJDBC4_2.5.18.1050.zip.
- For IBM DB2: the db2jcc.jar and db2jcc_license_cu.jar files, which are available from IBM
- For MySQL: a file with a name similar to mysql-connector-java-5.0.4-bin.jar, in a directory under the one into which you unzipped the download for the MySQL driver
- For Microsoft SQL Server or Sybase Adaptive Server: the jtds-n.n.n.jar that is included in the latest available JTDS .zip download.
- For PostgreSQL: postgresql-n.n-nnnn.jdbcn.jar, which is available from the Download page at <http://jdbc.postgresql.org/>
- For Teradata: tdgssconfig.jar and terajdbc4.jar, which are included (along with a readme.txt file) in the TeraJDBC__indep_indep.12.00.00.110.zip or TeraJDBC__indep_indep.12.00.00.110.tar download

Note:

As an alternative to using this preference, you can click **Help**, then **Check for Updates** to install the JTDS JDBC Driver for Microsoft SQL Server and the MySQL JDBE Driver as extensions.

To find a specific JDBC driver, see the appropriate website (for example, <http://www.mysql.com> for the MySQL Connector/J JDBC driver for MySQL, <http://sourceforge.net/projects/jtds/> for the jTDS driver for Microsoft SQL Server and Sybase Adaptive Server, <http://jdbc.postgresql.org/> for PostgreSQL, or search at <http://www.teradata.com/> for the JDBC driver for Teradata). For MySQL, use the MySQL 5.0 driver, not 5.1 or later, with SQL Developer release 1.5.

You must specify a third-party JDBC driver or install a driver using the Check for Updates feature before you can create a database connection to a third-party database of that associated type.

Database: User Defined Extensions

The User Defined Extensions pane specifies user-defined extensions that have been added. You can use this pane to add extensions that are not available through the Check for Updates feature. These extensions can be for user-defined reports, actions, editors, and navigators.

One use of the Database: User-Defined Extensions pane is to create a Shared Reports folder and to include an exported report under that folder: click Add Row, specify Type as REPORT, and for Location specify the XML file containing the exported report. The next time you restart SQL Developer, the Reports navigator will have a Shared Reports folder containing that report.

For more information about creating user-defined extensions, see:

- *How To create an XML User Defined Extension:* <https://wikis.oracle.com/display/oraclesqldeveloper/How+To+Create+an+XML+User+Defined+Extension>
- *Creating User Defined Extensions in Oracle SQL Developer:* http://www.oracle.com/webfolder/technetwork/tutorials/obe/db/11g/r2/prod/appdev/sqldev/srccodexmltext/xmltext_otn.htm

Database: Utilities

The Utilities pane specifies options that affect the behavior of Database utilities, including Export (Database: Utilities: Export) and Import (Database: Utilities: Import), when they are invoked using SQL Developer.

Database: Utilities: Cart

The Cart pane specifies the default locations for saving and opening Cart .xml files (**Default Cart Directory**) and for specifying .sql script files (**Open Script Directory**) to be executed before and after the generated master deployment script.

Database: Utilities: Cart: Cart Deploy

The Cart: Deploy pane specifies options that affect the behavior of the Cart for a deployment.

For example, if you do not want object names prefixed with the owner schema name in DDL creation statements (for example, if a generated script should create a BOOKS table and not a SMITH.BOOKS table), *deselect* (uncheck) the **Show Schema** option.

Default Cart Directory: Default location for the .zip file that contains the generated script files (including the master script file) for a deployment.

Database: Utilities: Difference

The Difference pane specifies options that affect the behavior of the Database Differences Wizard.

Database: Utilities: Export

The Export pane determines the default values used for the Database Export (Unload Database Objects and Data) wizard and for some other interfaces.

See also the panes for Database: Utilities: Export: Formats (CSV, Delimited, Excel, Fixed, HTML, PDF, SQL*Loader, Text, XML).

Export/View DDL Options: If this option is checked, the data definition language (DDL) statements for the database objects to be exported are included in the output

file, and the other options in this group affect the content and format of the DDL statements.

Note:

DDL is always exported in text format as a .sql file. DDL is not exported in other formats such as PDF or HTML.

Pretty Print: If this option is checked, the statements are attractively formatted in the output file, and the size of the file will be larger than it would otherwise be.

Terminator: If this option is checked, a line terminator character is inserted at the end of each line.

Show Schema: If this option is checked, the schema name is included in CREATE statements. If this option is not checked, the schema name is not included in CREATE and INSERT statements, which is convenient if you want to re-create the exported objects and data under a schema that has a different name.

Include Dependents: If this option is checked, objects that are dependent on the objects specified for export are also exported. For nonprivileged users, only dependent objects in their schema are exported; for privileged users, all dependent objects are exported.

Include BYTE Keyword: If this option is checked, column length specifications refer to bytes; if this option is not checked, column length specifications refer to characters.

Add Force to Views: If this option is checked, the FORCE keyword is added to any CREATE VIEW statements (resulting in CREATE OR REPLACE FORCE VIEW...) in the generated DDL during a database export operation. When the script is run later, the FORCE keyword causes the view to be created regardless of whether the base tables of the view or the referenced object types exist or the owner of the schema containing the view has privileges on them.

Include Grants: If this option is checked, GRANT statements are included for any grant objects on the exported objects. (However, grants on objects owned by the SYS schema are never exported.)

Include Drop Statement: If this option is checked, a DROP statement is included before each CREATE statement, to delete any existing objects with the same names. However, you may want to uncheck this option, and create a separate drop script that can be run to remove an older version of your objects before creation. This avoids the chance of accidentally removing an object you did not intend to drop.

Cascade Drops: If this option is checked, the DROP statements include the CASCADE keyword to cause dependent objects to be deleted also.

Storage: If this option is checked, any STORAGE clauses in definitions of the database objects are preserved in the exported DDL statements. If you do not want to use the current storage definitions (for example, if you will re-create the objects in a different system environment), uncheck this option.

Export Data: If this option is checked, the output file or files contain appropriate statements or data for inserting the data for an exported table or view; the specific output format is determined by the **Export Data: Format** setting. If this option is not checked, the data for an exported table or view is not exported; that is, only the DDL statements are included. If you check Export Data, all data in all tables in the selected schema is exported, unless you use the Filter Data tab to limit the data to be migrated.

Export Data: Format: Specifies the format to be used for exporting data for a table or view. For example, `insert` causes SQL INSERT statements to be included to insert the data, `loader` causes SQL*Loader files to be created, and `xls` causes a Microsoft Excel .xls file to be created.

Save As: Specifies how the output of an export operation is to be generated. You can also click **Browse** to find and select the location for the output:

- **Single File:** A single file contains both DDL and data.
- **Separate Files:** Each object is saved to a separate file in the specified directory.
- **Type Files:** Objects of the same type are saved to a single file in the specified directory.
- **Separate Directories:** A directory for each object type being exported is created. Files are created in the appropriate directory.
- **Worksheet:** Statements are sent to a SQL Worksheet window.
- **Clipboard:** Statements are copied to the clipboard.

Encoding: Character set to be used for encoding of the output file or files.

Open SQL File When Exported: If this option is checked, the output file is opened automatically when the export operation is complete.

Generate Controlling Script for Multiple SQL Files: If this option is checked, then if statements are generated in multiple .sql script files, a controlling (or master) .sql script file is also generated to run all the individual script files.

Database: Utilities: Export: Formats (CSV, Delimited, Excel, Fixed, HTML, PDF, SQL*Loader, Text, XML)

A separate pane is provided for specifying default attributes for each supported format for exported data. The displayed and editable attributes depend on the specific format.

Right Enclosure in Data is Doubled (CSV, Delimited, and Text formats) If this option is checked, then for CSV, delimited, and text format, if the right enclosure occurs in the data, it is doubled in the exported data file.

The **PDF Format** options include panes for options specific to cell, column, and table layout, headers and footers, and security. Note that if a table has multiple columns and the columns do not fit horizontally on a PDF page, then the Table Layout Horizontal Alignment option does not apply, but instead the table is split vertically spanning multiple pages and the information on each page is left-aligned.

PDF Format: Column Layout: Binary Large Object Column specifies whether to include BLOBs in the PDF document, to exclude BLOBs from the PDF document (and have the BLOB column be blank), or to create the BLOBs as separate files (in a *filename_blob* subdirectory) and create relative links in the PDF file to the BLOB files. (In this case, if you copy the PDF file, you must also remember to copy the *filename_blob* subdirectory and its files.) For the relative links option, you have the option to create instead only a .zip file that contains the PDF file and BLOB files.

If a pane includes **Header**, this option controls whether the first row is a header row or the first row of data.

Database: Utilities: Import

The Import pane determines the default values used for the Data Import Wizard (Load Data) wizard.

See also the panes for Database: Utilities: Import: File Formats (CSV, Delimited, Text) and Database: Utilities: Import: Import Methods (External Table, Insert, SQL Loader).

Default Import Directory: The default folder or directory for files containing data to be imported. You can click **Browse** to find and select the folder or directory.

File Encoding: Default character set used for encoding of the data to be imported.

Import Row Limit: If this option is checked, it specifies the maximum number of rows of data to be imported in a single import operation.

Preview Row Limit: If this option is checked, it specifies the maximum number of rows of data to be displayed in the preview pane before you proceed with the import operation. The lower the value, the faster the preview pane is populated. When the wizard creates a new table, the preview data is used to calculate the size of the columns; therefore, ensure that the preview is a good sample of the data.

Preview File Read Maximum: Specifies the maximum number of bytes of data to be displayed in the preview pane (if the specified preview row limit has not already been reached) before you proceed with the import operation. (You should set some maximum to ensure that very large files are not inadvertently read into memory.)

Database: Utilities: Import: File Formats (CSV, Delimited, Text)

A separate pane is provided for specifying default attributes for each supported format for imported data. The displayed and editable attributes depend on the specific format.

If a pane includes **Header**, this option controls whether the first row is a header row or the first row of data.

Delimiter: Identifies the character used to separate the data into columns. The delimiter character is not included in the data loaded. If the preview page does not show the data separated into columns, the correct delimiter is probably not specified. Examine the data in the preview area to determine the correct delimiter.

Line Terminator: Identifies the terminator for each line. The line terminator is not included in the data loaded. If the preview page shows the data in one single row, the correct terminator is not specified.

Left Enclosure and **Right Enclosure:** Enclosures are used for character data and are optional. Enclosures are not included in the data loaded.

Right Enclosure in Data is Doubled: If this option is checked, then for CSV, delimited, and text format, any right enclosure that is doubled in the data is loaded (imported) as a single right enclosure; also, the data preview will show such data as it will be loaded (that is, double right enclosures will be displayed as single right enclosures in the preview if this option is checked).

Database: Utilities: Import: Import Methods (External Table, Insert, SQL Loader)

The Import Methods pane lets you specify the default mode for importing data in the following situations:

- **Existing Table Import Method:** For importing data into an existing table, specifies the default method to be used for loading data.
- **New Table Import Method:** For creating a new table to hold the data to be imported, specifies the default method to be used for loading data.

A separate pane is provided for specifying default attributes for each supported method or tool to be used for exporting the data. The displayed and editable attributes depend on the specific method.

External Table Method preferences apply to both the External Table method and the Staging External Table method.

Identify Oracle directories for default, log, bad, and discard directories to be used as defaults for the data load. If any of the optional directories are not specified here or in the wizard, the default directory is used. Note these are Oracle directory objects that must be created before the external table can be used, and the appropriate permissions must be granted.

Staging Table Suffix: For Staging External Table, the suffix will be appended to the name of the table to create the name of the external table used as the staging table.

Commit and Drop Staging Table: Indicates the default setting for the commit and drop staging table option. If this option is enabled, the staging table will be dropped and the data will be committed at the end of the load.

Database: Worksheet

The Worksheet pane specifies options that affect the behavior of the SQL Worksheet.

Open a worksheet on connect: If this option is checked, a SQL Worksheet window for the connection is automatically opened when you open a database connection. If this option is not checked, you must use the Open SQL Worksheet right-click command or toolbar icon to open a SQL Worksheet.

New Worksheet to use unshared connection: If this option is checked, a separate unshared connection to the database is used for each new SQL Worksheet window that connects to a given database. If this option is not checked, the existing connection to the database is used for subsequent SQL Worksheet windows that connect to the database.

Close all worksheets on disconnect: If this option is checked, all SQL Worksheet windows for any database connection are automatically closed when that connection is disconnected.

Prompt for Save File on Close: If this option is checked, you are prompted to save changes when you close a SQL Worksheet if it contains any unsaved changes. If this option is not checked, any unsaved changes are discarded.

Grid in checker board or zebra pattern: If this option is checked, two different light-colored backgrounds are used for alternating rows of grid displays, to provide some visual contrast. If this option is not checked, all grid rows have a white background.

Max rows to print in a script: Limits the number of rows displayed.

Max lines in script output: Limits the number of lines output.

SQL History Limit: Maximum number of statements that can be stored in SQL History. Must be greater than 0 (zero). If you enter an invalid value, no value is stored in this field.

Default path to look for scripts: The default directory where SQL Developer looks when you run a script (using @). If you type a path, you can specify multiple delimited locations; if you click Browse, you can select a single location. In addition to any path that you specify, SQL Developer looks in the location specified by the SQLPATH environment variable.

Save bind variables to disk on exit: If this option is checked, bind variables that you enter when running a script are saved on disk for reuse. If you do not want bind

variable values stored on disk (for security or other reasons), be sure not to check this option.

Show query results in new tabs: If this option is checked, each query result tab showing the result of Run Statement (Ctrl-Enter) is "pinned" and a new tab is created for each new query's results. If this option is not checked, the result tab is reused for all queries. (This option does not apply to script output tabs produced by Run Script (F5). If you want to "pin" those, you must do so manually.)

Re-initialize on script exit command: If this option is checked, Script Runner context settings, such as a WHENEVER SQLERROR directive and substitution variables, are cleared when the script exits. If this option is not checked, these context settings are retained. For example, uncheck this option if you want to keep any WHENEVER SQLERROR setting specified in a script after the script exits.

Related Topics:

[SQL History](#) (page 1-47)

1.20.6 Debugger

The Debugger pane contains general options for the SQL Developer debugger. Other panes contain additional specific kinds of debugger options.

Show Tool Tip in Code Editor While Debugging: While debugging, if you stop at a breakpoint (or after stepping or pressing pause), you will be able to see a tooltip in the Code Editor when you hover the mouse over a variable name. The tooltip displays the current value of the variable.

Show Action Buttons in Log Window While Debugging: Displays the Actions list in the Log window toolbar while debugging. The Actions options allow you to wrap log messages, clear the log, and save the log in a text file.

Connection Retry Setting: Number of times that the debugger retries a connection to the debuggee process if the connection fails. If you see the message "Debugger unable to connect to local process" in the Log page after the debugger has tried many times to connect, review the remote debugger settings and the startup options of the debugger user.

Prompt for Debugger Host for Database Debugging: Select this option if JDeveloper is inside another network, such as a VPN network, and you are performing remote debugging to enable JDeveloper to request the appropriate IP address. When the JDeveloper is inside another network, the debugger normally sees the local IP address, but the database needs to use the VPN address to connect back. The option allows you to specify an alternate IP address for the database to use when establishing a connection back to the debugger.

If you do not select this option, the next time you perform local debugging you may see errors similar to these:

```
Executing PL/SQL: ALTER SESSION SET PLSQL_DEBUG=TRUE
Executing PL/SQL: CALL DBMS_DEBUG_JDWP.CONNECT_TCP( '172.17.0.5', '1841' )
XXX-NNNNN: failure establishing connection to debugger
XXX-NNNNN: TNS:protocol adapter error
```

Attempt to Break Method Evaluation Deadlocks: Attempts to break deadlocks formed during method evaluation.

Debugging Port Range: Sets the range of the port numbers used for remote debugging: **Minimum** (lowest allowed port number) and **Maximum** (highest allowed port number).

Start Debugging Option: Choose the option which describes how you want to start the debugger:

- **Run Until a Breakpoint Occurs:** When you start debugging, the debugger will let the program you are debugging execute until a breakpoint is reached.
- **Step Over:** When you start debugging, the debugger will let the program you are debugging execute until a method in a tracing-enabled class is reached, but it will not stop in class static initializer method.
- **Step Into:** When you start debugging, the debugger will let the program you are debugging execute until a method in a tracing-enabled class is reached.

Enable Change Tracking: You can control whether changes made during debugging are tracked, and the background and foreground colors to be used for changed items.

Debugger: Breakpoints

The Breakpoints pane sets the columns to appear in the Breakpoints pane and the scope of each breakpoint.

Debugger: Breakpoints: Default Actions

The Breakpoints: Default Actions pane sets defaults for actions to occur at breakpoints. These actions are the same as on the Actions tab in the Create/Edit Breakpoint dialog box.

Debugger: Data

The Data pane enables you to control the columns to appear in the debugger Data pane and aspects of how the data is displayed.

Debugger: Inspector

The Inspector pane enables you to control the columns to appear in the debugger Inspector pane and aspects of how the data is displayed.

Debugger: Smart Data

The Smart Data pane enables you to control the columns to appear in the debugger Smart Data pane and aspects of how the data is displayed.

Debugger: Stack

The Stack pane enables you to control the columns to appear in the debugger Stack pane and other options.

Debugger: ToolTip

The ToolTip pane enables you to control the columns to appear in the debugger ToolTip pane.

Debugger: Watches

The Watches pane enables you to control the columns to appear in the debugger Watches pane and aspects of how the data is displayed.

1.20.7 External Editor

The External Editor pane determines which external editor is called by SQL Developer when you try to edit binary large object (BLOB) data, such as image files, video files,

and other files created by certain applications. For each combination of MIME type and file extension, you can specify the executable application to be used to open and edit associated files.

MIME Type: MIME type of the data.

File Extension: File extension for files that contain BLOB data and that are associated with the MIME type.associated

Editor Location: Path to the editor to be used to open and edit files associated with this MIME type and file extension. To edit an existing path or to specify one if the cell is empty, click in the cell, and either modify the existing text or click Browse to find and select the executable file for the editor.

1.20.8 File Types

The File Types pane determines which file types and extensions will be opened by default by SQL Developer. The display shows each file extension, the associated file type, and a check mark if files with that extension are to be opened by SQL Developer by default, such as when a user double-clicks the file name.

Details area at bottom: You can modify the file type, content type (text or binary), and whether to open files with this extension automatically by SQL Developer.

To have files with a specific extension be opened by default by SQL Developer, click the file extension in the list, then check Open with SQL Developer in the Details area. This overrides any previous application association that may have been in effect for that file extension.

To add a file extension, click **Add** and specify the file extension (including the period). After adding the extension, you can modify its associated information by selecting it and using the Details area.

1.20.9 Global Ignore List

The Global Ignore List pane specifies filters that determine which files and file types will not be used in any processing.

New Filter: A file name or file type that you want to add to the list of files and file types (in the Filter box) that SQL Developer will ignore during all processing (if the filter is enabled, or checked). You can exclude a particular file by entering its complete file name, such as `mumble.txt`, or you can exclude all files of the same type by entering a construct that describes the file type, such as `*.txt`.

Add: Adds the new filter to the list in the Filter box.

Remove: Deletes the selected filter from the list in the Filter box.

Restore Defaults: Restores the contents of the Filter box to the SQL Developer defaults.

Filter: Contains the list of files and file types. For each item, if it is enabled (checked), the filter is enforced and the file or file type is ignored by SQL Developer; but if it is disabled (unchecked), the filter is not enforced.

1.20.10 Migration

The Migration pane contains options that affect the behavior of SQL Developer when you migrate schema objects and data from third-party databases to an Oracle database.

Default Repository: Migration repository to be used for storing the captured models and converted models.

Migration: Data Move Options

The Data Move Options pane contains options that affect the behavior when you migrate data from third-party databases to Oracle Database tables generated by the migration. This pane includes options that can be used for online data migration for MySQL and all supported third-party databases, and for offline data migration for MySQL, SQL Server, and Sybase Adaptive Server.

Oracle Representation for Zero Length String: The value to which Oracle converts zero-length strings in the source data. Can be a space (' ') or a null value (NULL).
Specific notes:

- For Sybase offline migrations, " is considered the same as a space (' ').
- For MySQL offline migrations, a null value is exported as 'NULL', which is handled as type VARCHAR2. You can specify another escape character by using the `--fields-escaped-by` option with the `mysqldump` command (for example, specifying `\N` for null or `\\` for `\`).

For MySQL offline migrations, the data is exported to a file named *table-name.txt*; so if you are moving data from two or more tables with the same name but in different schemas, rename files as needed so that they are all unique, and modify the `SQL*Loader .ctl` file accordingly.

Online: Options for online data move operations.

Number of Parallel Data Move Streams (online data moves): The number of internal connections created for simultaneous movement of data from the source database to the Oracle tables. Higher values may shorten the total time required, but will use more database resources during that time.

Number of Rows to Commit After (online data moves): During the data move operation, Oracle pauses to perform an automatic internal commit operation after each number of rows that you specify are moved from the source database to Oracle tables.

Lower values will cause a successful move operation to take more time; but if a failure occurs, it is likely that more source records will exist in the Oracle tables and that if the move operation is resumed, fewer source records will need to be moved. Higher values will cause a successful move operation to take less time; but if a failure occurs, it is likely that fewer source records will exist in the Oracle tables and that if the move operation is resumed, more source records will need to be moved.

Offline: Options for offline data move operations

End of Column Delimiter (offline data moves): String to indicate end of column.

End of Row Delimiter (offline data moves): String to indicate end of row.

Generic Date Mask (offline data moves): Format mask for dates, unless overridden by user-defined custom preferences.

Generic Timestamp Mask (offline data moves): Format mask for timestamps, unless overridden by user-defined custom preferences.

User-Defined Custom Preferences by Source Type (offline data moves): Lets you specify, for one or more source data types, a custom mapping for the function and format mask. Add one row for each mapping. For example, the following rows specify the Source Type, Function, and Mask for custom mappings for the Sybase data types `datetime`, `smalldatetime`, and `time`:

datetime	TO_TIMESTAMP	mon dd yyyy hh:mi:ss:ff3am
smalldatetime	TO_DATE	Mon dd yyyy hh:mi:am
time	TO_TIMESTAMP	hh:mi:ss:ff3am

Migration: Generation Options

The Generation Options pane contains options that are used when generating .sql script files for creating migrated database objects in the target schema.

One Single File or A File per Object Type or A file per Object: Determines how many files are created and their relative sizes. Having more files created might be less convenient, but may allow more flexibility with complex migration scenarios.

Generate Comments: Generates comments in the Oracle SQL statements.

Least Privilege Schema Migration: For migrating schema objects in a converted model to Oracle, causes CREATE USER, GRANT, and CONNECT statements *not* to be generated in the output scripts. You must then ensure that the scripts are run using a connection with sufficient privileges. You can select this option if the database user and connection that you want to use to run the scripts already exist, or if you plan to create them.

Generate Failed Objects: Causes objects that failed to be converted to be included in the generation script, so that you can make any desired changes and then run the script. If this option is not checked, objects that failed to be converted are not included in the generation script.

Generate Stored Procedure for Migrate Blobs Offline: Causes a stored procedure named CLOBtoBLOB_sqldeveloper (with execute access granted to public) to be created if the schema contains a BLOB (binary large object); this procedure is automatically called if you perform an offline capture. If this option is not checked, you will need to use the manual workaround described in [Populating the Destination Database Using the Data Files](#) (page 2-21). (After the offline capture, you can delete the CLOBtoBLOB_sqldeveloper procedure or remove execute access from public.)

Create Users: Causes separate database users to be created for each user specified for the objects to be migrated.

Generate Separate Emulation Package: Causes a separate database user named EMULATION to be created. The emulation package is created in the new EMULATION schema and is referenced by all other migrated users. If this option is not checked, no separate EMULATION user is created, and the emulation package is created within each migrated user. Generating a separate emulation user is usually the best practice because the emulation package is defined in one place, rather than having multiple copies of it. However, if you prefer each migrated user to be standalone and not need to reference anything from another user, then uncheck this option.

(Sybase) To Index-Organized Tables: Controls whether Sybase clustered unique constraints or clustered primary keys are converted to Oracle index-organized tables. NONE causes neither to be converted to index-organized tables (they are converted to Oracle unique constraints and primary keys, respectively). From Clustered Unique Constraints converts each clustered constraint into a primary key and creates an index-organized table if there is no primary key already. From Clustered Primary Keys creates an index-organized table if the primary key is clustered (the Sybase default).

Target Oracle Database 12c: Allows the migrated objects in the target Oracle Database schema to use features that are new in Oracle Database 12c (Release 12.1). If this option is not checked, features that are new in Oracle Database 12c will not be considered for

inclusion. Checking this option also adds `QUOTA UNLIMITED ON SYSTEM` to the `CREATE USER` statements in generated scripts. For example (for a user named `Emulation`):

```
CREATE USER Emulation IDENTIFIED BY Emulation DEFAULT TABLESPACE SYSTEM TEMPORARY
TABLESPACE TEMP QUOTA UNLIMITED ON SYSTEM;
GRANT CREATE SESSION, RESOURCE, CREATE VIEW, CREATE MATERIALIZED VIEW, CREATE
SYNONYM, CREATE PUBLIC SYNONYM, ALTER SESSION, CREATE SQL TRANSLATION PROFILE TO
Emulation;
```

(If you want to use a different tablespace than `SYSTEM`, modify the scripts accordingly.)

Object Types: Specifies the types of database objects to be included by default in migrations. You can toggle all or toggle individual types to select and deselect types.

Migration: Identifier Options

The Identifier Options pane contains options that apply to object identifiers during migrations.

Prepended to All Identifier Names (Microsoft SQL Server and Sybase Adaptive Server migrations only): A string to be added at the beginning of the name of migrated objects. For example, if you specify the string as `XYZ_`, and if a source table is named `EMPLOYEES`, the migrated table will be named `XYZ_EMPLOYEES`. (Be aware of any object name length restrictions if you use this option.)

Is Quoted Identifier On (Microsoft SQL Server and Sybase Adaptive Server migrations only): If this option is enabled, quotation marks (double-quotes) can be used to refer to identifiers (for example, `SELECT "Col 1" from "Table 1"`); if this option is not enabled, quotation marks identify string literals. **Important:** The setting of this option must match the setting in the source database to be migrated.

Migration: Teradata

The Teradata pane contains options that apply to migrations from Teradata.

Fast Export Min Sessions and **Fast Export Max Sessions:** The minimum and maximum number, respectively, of sessions to be used for the Teradata FastExport utility.

Migration: Translators

The Translators pane contains options that relate to conversion of stored procedures and functions from their source database format to Oracle format. (These options apply only to migrations from Microsoft SQL Server and Sybase Adaptive Server.)

Default Source Date Format: Default date format mask to be used when casting string literals to dates in stored procedures and functions.

Variable Name Prefix: String to be used as the prefix in the names of resulting variables.

In Parameter Prefix: String to be used as the prefix in the names of resulting input parameters.

Query Assignment Translation: Option to determine what is generated for a query assignment: only the assignment, assignment with exception handling logic, or assignment using a cursor `LOOP ... END LOOP` structure to fetch each row of the query into variables.

Display AST: If this option is checked, the abstract syntax tree (AST) is displayed in the Source Tree pane of the Translation Scratch Editor window if you perform a translation.

Generate Compound Triggers: If this option is checked, then depending on the source code, SQL Developer can convert a Sybase or SQL Server trigger to an Oracle compound trigger, which uses two temporary tables to replicate the inserted and deleted tables in Sybase and SQL Server. In such cases, this can enable the conversion of logic that cannot otherwise be converted.

Oracle Database 12c Features: For each of the listed features that are available only starting with Oracle Database Release 12.1 or later, you can allow or disallow it to be used (where relevant and appropriate) when SQL Developer creates the converted stored procedures and functions.

Related Topics:

[SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

[Creating Data Files From MySQL](#) (page 2-20)

1.20.11 Mouse Actions

The Mouse Actions pane specifies text to be displayed on hover-related mouse actions over relevant object names.

Popup Name: The type of information to be displayed: **Data Values** (value of the item under the mouse pointer, such as the value of a variable during a debugging operation), **Documentation** (documentation on the item under the mouse pointer, such as Javadoc on a method call), or **Source** (source code of the item under the mouse pointer, such as the source code of a method).

Activate Via: Use action with the mouse cursor to activate the display: Hover, or Hover while pressing one or two specified modifier keys.

Description: Description of the associated Popup Name entry.

Smart Enabled: If this option is checked, then the text for the relevant type of information is displayed if Smart Popup is also checked.

Smart Popup: If this option is checked, the relevant text for the first smart-enabled popup is displayed for the item under the mouse pointer.

1.20.12 Shortcut Keys (Accelerator Keys)

The Shortcut Keys pane enables you to view and customize the shortcut key (also called accelerator key) mappings for SQL Developer.

Show Mapped Commands Only: If this option is checked, only shortcut keys with mappings are displayed.

More Actions:

- **Export:** Exports the shortcut key definitions to an XML file.
- **Import:** Imports the shortcut key definitions from a previously exported XML file.
- **Load Keyboard Scheme:** Drops all current shortcut key mappings and sets the mappings in the specified scheme. (This option was called Load Preset in some previous releases.) *If you have made changes to the mappings and want to restore the default settings, select **Default**.*

Category: Select All or a specific category (Code Editor, Database, Debug, Edit, and so on), to control which actions are displayed.

Command: The actions for the selected category. When you select an action, any existing shortcut key mappings are displayed.

Shortcut: Any existing key mappings for the selected action. To remove an existing key mapping, select it and click Remove.

New Shortcut: The new shortcut key to be associated with the action. Press and hold the desired modifier key, then press the other key. For example, to associate Ctrl+J with an action, press and hold the Ctrl key, then press the j key. If any actions are currently associated with that shortcut key, they are listed in the Current Assignment box.

Conflicts: A read-only display of the current action, if any, that is mapped to the shortcut key that you specified in the New Shortcut box.

1.20.13 SSH (Secure Shell)

SSH preferences are related to creating SSH (Secure Shell) connections.

Use Known Hosts File: If this option is checked, specify the file of known hosts to be used.

Related Topics:

[Connections with SSH Authentication](#) (page 1-33)

1.20.14 Unit Test Parameters

Unit Test Parameters preferences affect the behavior of the SQL Developer unit testing feature.

Configuration set to use for lookups: Lookup category to be used for automatically generating test implementations when you create a unit test. The list includes the default category and any user-added categories.

Related Topics:

[Automatically Creating Implementations](#) (page 4-9)

1.20.15 Usage Reporting

In SQL Developer, Data Modeler, and some other applications, the Usage Reporting user preference and a related dialog box ask for your consent to Oracle usage reporting. If you consent, automated reports can occasionally be sent to Oracle describing the product features in use. No personally identifiable information will be sent and the report will not affect performance. You can review Oracle's privacy policy by clicking the **privacy policy** link.

Allow automated usage reporting to Oracle: Determines whether you consent to usage reporting.

1.20.16 Versioning

Versioning preferences affect the behavior of the version control and management systems that you have available for use with SQL Developer.

- [Versioning: Git](#) (page 1-117)

- [Versioning: Subversion](#) (page 1-117)
- [Versioning: CVS](#) (page 1-118)
- [Versioning: Perforce](#) (page 1-120)

1.20.16.1 Versioning: Git

The Versioning: Git pane introduces options for use with the Git version control system.

Versioning: Git: General

The Git: General pane specifies environment settings.

Use Navigator State Overlay Icons: If this option is enabled, state overlay icons are used. State overlay icons are small symbols associated with object names in the navigators. They indicate the state of version-controlled files (for example, "up to date").

Use Navigator State Overlay Labels: If this option is enabled, state overlay labels are used. State overlay labels are tooltips associated with object names in the navigators.

Automatically Add New Files on Committing Working Tree: If this option is enabled, any new files you have created in your working copy are automatically added to the Git repository whenever you commit any individual file. Otherwise, Git will not add new files when you commit changes; you must continue to add new files to Git explicitly.

Write Messages to Log Window: If this option is enabled, Git messages are written to the Messages - Log window. (If that window is not visible, click View > Log to display it.)

Versioning: Git: Version Tools

The Git: Version Tools pane specifies options for the Outgoing Changes Commit and Add dialog box.

Use Outgoing Changes Commit and Add Dialog: Enables you to make optimum use of limited screen space when the Pending Changes window is open. You can save screen space by not showing the Comments area of the Pending Changes window, but you might still want to add comments before a commit action. You can choose the circumstances under which the Outgoing Changes Commit and Add dialog is opened: always, only when the Comments area of the Pending Changes window is hidden, or never.

1.20.16.2 Versioning: Subversion

The Subversion pane specifies the Subversion client to use with SQL Developer.

Versioning: Comment Templates

The Comment Templates pane specifies templates for comments to be used with commit operations. For example, a template might contain text like the following:

```
Problem Description (with bug ID if any):
Fix Description:
```

You can add, edit, and remove comment templates, and you can export templates to an XML file or import templates that had previously been exported.

Versioning: Subversion: General

The Subversion: General pane specifies environment settings and the operation timeout.

Use Navigator State Overlay Icons: If this option is enabled, state overlay icons are used. State overlay icons are small symbols associated with object names in the navigators. They indicate the state of version-controlled files (for example, "up to date").

Use Navigator State Overlay Labels: If this option is enabled, state overlay labels are used. State overlay labels are tooltips associated with object names in the navigators.

Automatically Add New Files on Committing Working Copy: If this option is enabled, any new files you have created in your working copy are automatically added to the Subversion repository whenever you commit any individual file. Otherwise, Subversion will not add new files when you commit changes; you must continue to add new files to Subversion explicitly.

Automatically Lock Files with svn:needs-lock Property After Checkout: If this option is enabled, files you check out from the repository are automatically locked, preventing other team members from checking them out until you release the files.

Use Merge Wizard for Subversion Merging: If this option is enabled, the Merge wizard rather than the Merge dialog box is invoked for merge requests.

Operation Timeout: Maximum number of seconds, minutes, or hours allowed for Subversion operations to complete.

Edit Subversion Configuration File: To modify the Subversion file directly, click **Edit "server"**. You can make changes to server-specific protocol parameters such as proxy host, proxy port, timeout, compression, and other values. Lines beginning with # are interpreted as comments.

Versioning: Subversion: Version Tools

The Subversion: Version Tools pane specifies options for the pending changes window and the merge editor.

Use Outgoing Changes Commit Dialog: Enables you to make optimum use of limited screen space when the Pending Changes window is open. You can save screen space by not showing the Comments area of the Pending Changes window, but you might still want to add comments before a commit action. You can choose the circumstances under which the Commit dialog is opened: always, only when the Comments area of the Pending Changes window is hidden, or never.

Incoming Changes Timer Interval: The frequency at which the change status of files is checked.

Merge Editor: Specifies whether files are merged locally or at the server.

1.20.16.3 Versioning: CVS

(Available only if you added support for CVS.)

The CVS pane specifies options for use with CVS (Concurrent Versions System).

CVS Client: Internal to Oracle SQL Developer (installed with SQL Developer) or **External Executable** (separately installed CVS client, for which you must specify the name or path).

- **Name on System Path:** Name of the CVS server executable. The default (cvs) is correct for most installations. This option assumes that the name of the CVS server executable is on the system path.

- **Path from Environment:** Location of the CVS server executable, especially if there is more than one on the system path. The selection area will list all instances of the CVS server executable known to the local system. You may have more than one version of CVS installed: this option lets you specify which of them to use with SQL Developer.
- **Other Path:** Location of the CVS server executable, if it is not on the system path at all.

Run CVS in Edit/Watch Mode: If this option is enabled, you coordinate access to files by declaring an editor for them through CVS, after which they may be modified. Only those files that you check out after changing this preference will be affected. If this option is disabled, the edit and watch commands on the Team menu are disabled.

State Overlay Scheme: Scheme for the icons displayed alongside folder and file names in the navigators to indicate their versioning status.

Versioning: CVS: Commands

The CVS: Commands pane sets options for CVS source control. Some options are not available when using the internal CVS client.

Enable Advanced Controls: If this option is enabled, advanced CVS controls are shown in dialog boxes. If you find that you use only basic CVS features, you might want to use SQL Developer without advanced controls, to reduce complexity and save screen space.

Global Options: Run Quietly: If this option is enabled, informational messages are suppressed.

Global Options: Do not Log Commands: If this option is enabled, CVS commands are not logged in the repository command history.

Global Options: Encrypt: If this option is enabled, all communication between the client and the server is encrypted. Encryption support is not available in CVS by default; it must be enabled using a special configuration option when you build CVS.

Set Compression Level (z): If this option is enabled, you can set the compression level for files sent between client and server. The level can be set from Minimum (high speed, low compression) to Maximum (low speed, high compression).

Keyword Substitution Mode: CVS uses keyword substitution modes to insert revision information into files when they are checked out or updated. This option controls the mode of replacement for keyword substitution in versioned files:

- **Automatic:** The default, recommended option.
- **Keyword-Only Mode:** Generates only keyword names in keyword strings and omits their values. This option is useful for disregarding differences due to keyword substitution when comparing different revisions of a file.
- **Keyword-Value Mode:** Generates keyword strings using the default form.
- **Keyword-Value-Locker Mode:** Like the keyword-value mode, except that the name of the locker is always inserted if the given revision is currently locked.
- **Old-Contents Mode:** Generates the old keyword string, present in the working file just before it was checked in.
- **Value-Only Mode:** Generates only keyword values for keyword strings. This can help generate files in programming languages where it is hard to strip keyword delimiters from a string. However, further keyword substitution cannot be

performed once the keyword names are removed, so this option should be used with care.

On Commit: Use Comment Templates: If this option is enabled, your commit comments will be entered through template forms. The forms are set up by the CVS system administrator. There may be different forms for different circumstances and installations, and it may be that none of them are suitable for your commit comments. In this case, this preference lets you disable the use of all forms.

On Commit: Automatically Add Files: If this option is enabled, local files are added to the CVS repository whenever you perform a commit action.

Create Backup Files on Remove: If this option is enabled, backup copies are made of files that are removed through actions of the source control system.

Versioning: CVS: General

The CVS: General pane specifies environment settings and the operation timeout.

Use Navigator State Overlay Icons: If this option is enabled, state overlay icons are used. State overlay icons are small symbols associated with object names in the navigators. They indicate the state of version-controlled files (for example, "up to date").

Use Navigator State Overlay Labels: If this option is enabled, state overlay labels are used. State overlay labels are tooltips associated with object names in the navigators.

Automatically Make Files Editable: If this option is enabled, an editor is automatically used on a data file when you start to change it. (If you edit a file unintentionally, immediately use Versioning, then Unedit to revert.)

Operation Timeout: Maximum time allowed for CVS operations to complete.

Versioning: CVS: Navigator Labels

The CVS: Navigator Labels pane specifies formatting for CVS information appears on navigator nodes and tool tips. For a full explanation of keyword substitution modes, see the CVS documentation.

Versioning: CVS: Version Tools

The CVS: Version Tools pane specifies options for the pending changes window and the merge editor.

Use Outgoing Changes Commit Dialog: Enables you to make optimum use of limited screen space when the Pending Changes window is open. You can save screen space by not showing the Comments area of the Pending Changes window, but you might still want to add comments before a commit action. You can choose the circumstances under which the Commit dialog is opened: always, only when the Comments area of the Pending Changes window is hidden, or never.

Incoming Changes Timer Interval: The frequency at which the change status of files is checked.

Merge Editor: Specifies whether files are merged locally or at the server.

1.20.16.4 Versioning: Perforce

(Available only if you added support for Perforce.)

The Perforce pane lets you specify the Perforce client executable to invoke from SQL Developer. Specify the name on the system path, the full file path in your local environment, or another path.

Versioning: Comment Template

The Comment Templates pane specifies templates for comments to be used with commit operations. For example, a template might contain text like the following:

```
Problem Description (with bug ID if any):
Fix Description:
```

You can add, edit, and remove comment templates, and you can export templates to an XML file or import templates that had previously been exported.

Versioning: Perforce: General

The Perforce: General pane specifies environment settings and the operation timeout.

Use Navigator State Overlay Icons: If this option is enabled, state overlay icons are used. State overlay icons are small symbols associated with object names in the navigators. They indicate the state of version-controlled files (for example, "up to date").

Use Navigator State Overlay Labels: If this option is enabled, state overlay labels are used. State overlay labels are tooltips associated with object names in the navigators.

Write Messages to Log Window: If this option is enabled, Perforce messages are written to the Messages - Log window. (If that window is not visible, click View > Log to display it.)

Automatically Open Files for Edit: If this option is enabled, data files are marked as open for edit, and then opened so that you can immediately start making changes. (If you edit a file unintentionally, immediately use VersioningRevert to revert.)

Operation Timeout: Maximum number of seconds, minutes, or hours allowed for Perforce operations to complete.

Versioning: Perforce: Version Tools

The Perforce: Version Tools pane specifies options for the pending changes window and the merge editor.

Use Outgoing Submit Dialog: Enables you to make optimum use of limited screen space when the Pending Changes window is open. You can save screen space by not showing the Comments area of the Pending Changes window, but you might still want to add comments before a commit action. You can choose the circumstances under which the Commit dialog is opened: always, only when the Comments area of the Pending Changes window is hidden, or never.

Merge Editor: Specifies whether files are merged locally or at the server.

1.20.17 Web Browser and Proxy

The Web Browser and Proxy settings are relevant only when you use the Check for Updates feature (click **Help**, then **Check for Updates**), and only if your system is behind a firewall.

Web Browsers

Displays the available web browsers and the default browser for Check for Update operations. You can click under Default to change the default browser.

For each browser, you can determine whether it is the default, and you can see and optionally change its name, the path to the application's executable file, application command parameters, and the icon.

Proxy Settings

You can choose no proxy, system default proxy settings, or manually specified proxy settings for Check for Update operations. For manually specified settings, check your Web browser options or preferences for the appropriate values for these fields.

Internet Files

You can choose whether to enable Internet cookies for Check for Update operations.

Clear All Cookies: Clears all existing cookies.

1.20.18 XML Schemas

The XML Schemas pane lets you view all the currently registered XML schemas, add new schemas to support additional namespaces and elements, remove existing schemas, and unload schemas from memory.

SQL Developer Schemas for XML Editing: Lists the names and locations of the pre-registered schemas currently available when editing XML documents and the file extension with which each schema is associated.

User Schemas for XML Editing: Lists the names and locations of the schemas you have added since installing SQL Developer that are available when editing XML documents, and the file extension with which each schema is associated.

Add: Displays the **Register Schema for File Type** dialog box for specifying a new schema to add to the list of User Schemas.

Remove: Removes the selected schema from the list.

Edit: Displays a dialog box in which you can modify a previously registered schema.

For adding or editing, you can specify the location (file system or URL) and the file extension to register the schema for a specific file type.

Clear Cache: Unloads all currently loaded schemas from memory when you modify a schema. Any needed schemas will then be reloaded, including the modified schema.

1.21 Location of User-Related Information

SQL Developer stores user-related information in several places, with the specific location depending on the operating system and certain environment specifications. User-related information includes user-defined reports, user-defined snippets, SQL Worksheet history, code templates, and SQL Developer user preferences. In most cases, your user-related information is stored outside the SQL Developer installation directory hierarchy, so that it is preserved if you delete that directory and install a new version.

The user-related information is stored in or under the IDE_USER_DIR environment variable location, if defined; otherwise as indicated in the following table, which shows the typical default locations (under a directory or in a file) for specific types of resources on different operating systems. (Note the period in the name of any directory named `.sqldeveloper`.)

Table 1-1 Default Locations for User-Related Information

Resource Type	System (Windows, Linux, or Mac OS X)
User-defined reports	Windows: C:\Documents and Settings\ <i><user-name></i> \Application Data\SQL Developer\UserReports.xml Linux or Mac OS X: ~/.sqldeveloper/UserReports.xml

Table 1-1 (Cont.) Default Locations for User-Related Information

Resource Type	System (Windows, Linux, or Mac OS X)
User-defined snippets	Windows: C:\Documents and Settings\ <i><user-name></i> \Application Data\SQL Developer\UserSnippets.xml Linux: ~/.sqldeveloper/UserSnippets.xml Mac OS X: /Users/<Your user>/Library/Application Support/SQLDeveloper/UserSnippets.xml
SQL history	Windows: C:\Documents and Settings\ <i><user-name></i> \Application Data\SQL Developer\SqlHistory.xml Linux: ~/.sqldeveloper/SqlHistory.xml Mac OS X: /Users/<Your user>/Library/Application Support/SQLDeveloper/SqlHistory.xml
Code templates	Windows: C:\Documents and Settings\ <i><user-name></i> \Application Data\SQL Developer\CodeTemplate.xml Linux: ~/.sqldeveloper/CodeTemplate.xml Mac OS X: /Users/<Your user>/Library/Application Support/SQLDeveloper/CodeTemplate.xml
SQL Developer user preferences	Windows: C:\Documents and Settings\ <i><user-name></i> \Application Data\SQL Developer\system <i>n.n.n.n.n</i> Linux or Mac OS X: ~/.sqldeveloper/system <i>n.n.n.n.n</i>

If you want to prevent other users from accessing your user-specific SQL Developer information, you must ensure that the appropriate permissions are set on the directory where that information is stored or on a directory above it in the path hierarchy. For example, on a Windows system you may want to ensure that the `SQL Developer` folder and the `<user-name>\Application Data\SQL Developer` folder under `Documents and Settings` are not sharable; and on a Linux or Mac OS X system you may want to ensure that the `~/.sqldeveloper` directory is not world-readable.

1.22 Java Development Kit (JDK) Location

SQL Developer stores the location of the Java Development Kit (JDK) that it is using in a configuration file named `product.conf`, which is located under the user directory (operating system-specific). For example, on a Windows system it might be in a location with the following format:

```
C:\Users\username\AppData\Roaming\sqldeveloper\i.n.n.n\product.conf
```

If you have made an error in your JDK specification or configuration, you can try editing the `SetJavaHome` line in `product.conf`.

1.23 Command-Line Interface for SQL Developer

For certain operations, you can invoke SQL Developer from the command line as an alternative to the graphical user interface. To use the command-line interface, go to the `sqldeveloper\sqldeveloper\bin` folder or `sqldeveloper/sqldeveloper/bin` directory under the location where you installed SQL Developer, and enter `sdcli`.

For example, if you installed SQL Developer on Windows under C:\, enter the following commands:

```
C:\>cd c:\sqldeveloper\sqldeveloper\bin
```

```
C:\sqldeveloper\sqldeveloper\bin>sdcli
Available features:
cart: Database Cart Batch Tasks
dba: Basic Batch DBA Tasks
format: SQL Format Task
migration: Database Migration Tasks
reports: Basic Batch Reporting Tasks
unittest: Unit Testing Batch Tasks
```

If you enter just `sdcli`, the list of available features shows each `sdcli` command and its associated feature. To get brief information about commands for a specific feature, enter `sdcli <feature-name>`. For example:

```
C:\sqldeveloper\sqldeveloper\bin>sdcli migration
HELP:
Syntax:
  Migration -help|-h=<actions> to get help on one or more actions.
where
  <actions>      could be one or more (comma separated list) actions or guide
  guide provides a walk through of a typical migration
Examples:
  Migration -help=capture
  Migration -h=capture,convert
  Migration -h=guide
```

```
Valid actions:
capture, convert, datamove, delcaptured, delconn, delconverted, driver, generate
, guide, idmap, info, init, lscaptured, lsconn, lsconverted, mkconn, qm, runsql,
scan and translate
```

If a command has a `-logfile` option, the generated log file is empty if the command failed; otherwise, it contains one or more descriptive messages about the operation.

1.24 Data Modeler in SQL Developer

SQL Developer includes an integrated version of SQL Developer Data Modeler. You can create, open, import, and save a database design, and you can create, modify, and delete Data Modeler objects.

To display Data Modeler in the SQL Developer UI, click **View**, then **Data Modeler**, then **Browser**.

- To create a new design, right-click the **Designs** node and select **New Design**.
- To open an existing design, click **File**, then **Data Modeler**, then **Open**.

If the design is one that you have recently opened, you can click **File**, then **Data Modeler**, then **Recent Designs** and select the design.

To access other Data Modeler features (such as specifying design rules and user preferences), click **Tools**, then **Data Modeler**.

The **Model** tab in a table display includes **Open in Data Modeler**, which enables you to open the table and sometimes related tables in a Data Modeler diagram view.

- If you right-click a table in the diagram and select **Show Parent and Child Tables**, any parent and child tables are added to the display (if they are not already included).
- If you right-click a table in the diagram and select **Properties**, a Data Modeler table properties display is opened in read-only mode.

For information about using Data Modeler:

- Use the online help: click **Help** or press the F1 key in the Data Modeler interface; or in the SQL Developer menus click **Help**, then **Table of Contents**, then look for the *Data Modeler Concepts and Usage* topic

Related Topics:

Oracle SQL Developer Data Modeler User's Guide

1.25 Oracle TimesTen In-Memory Database Support

When you connect to an Oracle TimesTen In-Memory Database, the available types of objects that you can work with include several that apply to an Oracle Database, and the following that are specific to TimesTen:

- Cache groups
- Replication schemes

To create a connection to a TimesTen database, use the TimesTen tab in the Create/Edit/Select Database Connection dialog box.

Related Topics:

Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide

1.26 Using the Help

SQL Developer provides a Help menu and context-sensitive help (click the Help button or press the F1 key in certain contexts). Much of the help content is also in *Oracle SQL Developer User's Guide*, which is in the SQL Developer Documentation Library.

Help is displayed in the Help Center window, which has a Contents pane on the left, a Search box at the top right, and a help topic display pane under the Search box. You can move the horizontal divider to change the pane sizes (for example, to make the Contents pane narrower, to allow more room for the help topic content). You can also resize and reposition the Help Center window.

For Search, you can click the icon (binoculars) to see search options: case sensitive (Match case) or case insensitive; and whether to match topics based on all specified words, any specified words, or a Boolean expression.

The **Keep on Top** button toggles whether the Help Center window is kept on top of the display when you switch focus (click) back in the SQL Developer window.

To print a help topic, display it in the topic display pane and click the Print icon at the top of the pane.

To increase or decrease the size of the font in the help topic viewer, click the Change Font Size (A) icon in the Help Center topic display area toolbar, then select Increase Font Size or Decrease Font Size. This setting is preserved only for the duration of the

current help pane or window; therefore, you may want to keep the Help Center window open after setting the help text font to your preferred size.

1.27 SQL Developer Usage Suggestions

This topic contains several usage suggestions for SQL Developer. (The material was displayed through the "Tip of the Day" feature for English locales in previous releases.)

1.27.1 SQL History Shortcuts

Using **Ctrl+up-arrow** or **Ctrl+down-arrow** in the SQL Worksheet replaces the contents of the SQL Worksheet with lines of code from the SQL History. You can step up and down through the SQL History.

To view the SQL History in the SQL Worksheet, press **F8** or click **View**, then **History**.

1.27.2 Unshared Worksheets

To create a separate unshared worksheet for a connection, press **Ctrl+Shift+N**.

1.27.3 SQL Worksheet Bookmarks

If you have many SQL Worksheets open, you can assign a bookmark number to each and then easily navigate among them. To create a bookmark, click the worksheet's tab and press **Alt+Shift+number** (for example, Alt+Shift+1) or click **Window**, then **Assign File Accelerator**, then the appropriate selection (for example, Assign to Alt+1). The number now appears as a small superscript in the tab.

To switch to a worksheet that has a bookmark, press **Alt+number** (for example, Alt+1).

1.27.4 Oracle Data Miner

Interested in Oracle Data Miner? Click **Help**, then **Data Mining**.

1.27.5 Formatted Display of SQL Trace (.trc) Files

To see a formatted display of a SQL Trace file, drag the *.trc file onto the area above the SQL Worksheet (or open it by clicking **File**, then **Open**).

1.27.6 Keyboard Navigation: Alt + Page Down/Up to Move Among Tabs

You can press **Alt + Page Down** to move to the **next tab**, and **Alt + Page Up** to move to the **previous tab**, in a tabbed editor or display window, such as with the tabs for Table and View grid displays (Columns, Data, and so on).

1.27.7 Folders for Organizing Connections

You can group connections into folders. Right-click a connection name and select **Add to Folder**.

Related Topics:

[Using Folders to Group Connections](#) (page 1-31)

1.27.8 Third-Party Databases and SQL Developer

In addition to Oracle databases, SQL Developer works with MySQL and several third-party databases, such as Sybase, Microsoft SQL Server, and IBM DB2.

1.27.9 Debugger Ports and Firewalls

The SQL Debugger by default uses ports 40000 to 49000. If you cannot get the debugger to start, make sure that you are not being blocked by a firewall on these ports.

1.27.10 Viewing Multiple Tables

You can have tabs open for more than one table. Just click the Freeze View button (it looks like a push pin) when you are viewing a table; and when you click to display another table, the tab for the first table will remain open.

1.27.11 Customizing SQL Developer Appearance

You can use the **Look and Feel** (platform) option under Environment preferences to customize the appearance of the SQL Developer window.

1.27.12 Maximizing Tab Panes

You can often maximize a display pane (such as a SQL Worksheet) by double-clicking its tab.

To restore the SQL Developer window to its original display, double-click the tab again.

1.27.13 Default Path for Running Scripts

You can set a default path for SQL Developer to use if you run a SQL script file (for example, *@my_script.sql*) without specifying the path. See the Database: Worksheet Parameters preferences.

1.27.14 Shutting Down and Restarting the Database

A user with SYSDBA privileges can shut down and restart the database from within SQL Developer, if a listener is running with a static listener configured for the database. Right-click the connection name and select **Manage Database**.

1.27.15 Feature Requests

Do you have a SQL Developer feature request? Log it at the Oracle Technology Network: go to <http://sqldeveloper.oracle.com> and select the **Feature Requests** link.

1.27.16 Discussion Forums

Would you like to share and search information, questions, and comments about SQL Developer and Data Miner? Visit our discussion forums:

SQL Developer: https://community.oracle.com/community/database/developer-tools/sql_developer

Data Mining: <http://forums.oracle.com/forums/forum.jspa?forumID=55>

1.27.17 Help Text Font Size

You can use a button in the Help Center window display pane (right side) to change the help text display size: click **Increase Font Size** or **Decrease Font Size** (repeatedly if necessary) until the size is right for you. (If the text appears blurry, try decreasing the size.)

Suggestion: Don't close the Help Center window, because font size changes will be in effect only for as long as the current window is open. Consider using the **Keep on Top** toggle in the Help Center window.

1.27.18 Procedure and Function Signatures

To see the signature (format, including parameters) of a procedure or function in a PL/SQL package, expand the package (under Packages in the Connections navigator), and place the mouse pointer over the procedure or function name.

1.27.19 Type-Ahead in Navigators

Many navigators that use a tree support type-ahead to find and open an object. For example, expand the Tables node under a connection and start typing a table name.

Note: This works only on nodes when the child nodes are visible. For example, if the Tables node is not expanded to display the individual tables, typing the name of a table will not find and open it.

1.27.20 Extended Paste

If you have cut or copied multiple things to the clipboard and want to paste something other than the most recent copy, you can use extended paste to display a dialog box to select which one to paste. Press **Ctrl+Shift+V**; or click **Edit**, then **Extended Paste**.

1.27.21 Closing Tabbed Windows Using the Mouse Wheel

To close a tabbed editor or display window, click its tab with the mouse wheel.

1.27.22 Go to Last Edit Location

If you have made edits in several editing windows and are now in a different window, and if you want to return to where you made the last edit, press **Ctrl+Shift+Backspace**; or click **Navigate**, then **Go to Last Edit**.

1.27.23 Closing Tabbed Windows Using the Context Menu

To close a tabbed editor or display window, right-click and select **Close** from the context menu.

1.27.24 List of All Open Windows

To see a list of all open tabbed windows, click the small button with the drop-down arrow, located to the right of the tabs and over the tabbed window vertical scroll bar.

To go to one of the listed windows, select it from the drop-down list.

1.27.25 Go to Subprogram Implementation from Package Window

In the window for a package definition, you can press **Ctrl+click** on a procedure or function name to perform the Open Declaration command, which opens the procedure or function implementation (body specification) in a new window.

1.27.26 Select Multiple Table or Column Names in Completion Insight

When entering or editing a SELECT query, you can select multiple tables and columns from the completion insight popup window. Aliases are provided for column and table names if the **Generate Column/Table Aliases Automatically** preference for Code Editor: Completion Insight is enabled.

1.27.27 Startup Time and Automatic Check for Updates

If the startup time for SQL Developer seems too slow, consider **disabling** the *Automatically Check for Updates* option (Tools, Preferences, Extensions).

If this option is enabled, it can increase the startup time. You can manually check for updates by clicking **Help**, then **Check for Updates**.

1.28 For More Information About SQL Developer

For more information about SQL Developer and related topics, you may find the following resources helpful:

- SQL Developer Start Page, which contains links for tutorials, online demonstrations, documentation, and other resources. This page has two tabs: Get Started and Community. (If the Start Page tab is not visible, click **Help**, then **Start Page**).
- SQL Developer home page (OTN), which includes links for downloads, white papers, tutorials, viewlets (demonstrations), blogs, a discussion forum, and other sources of information: <http://www.oracle.com/technetwork/developer-tools/sql-developer/>
- PL/SQL page on OTN: <http://www.oracle.com/technetwork/database/features/plsql/>
- Oracle Accessibility site: <http://www.oracle.com/accessibility/>
- Oracle Corporate site: <http://www.oracle.com/>

SQL Developer: Migrating Third-Party Databases

Migration is the process of copying the schema objects and data from a source MySQL or third-party (non-Oracle) database, such as Microsoft SQL Server, Sybase Adaptive Server, or IBM DB2 (UDB), to Oracle Database. You can perform the migration in an efficient, largely automated way.

Thus, you have two options for working with databases other than Oracle Database in SQL Developer:

- Creating database connections so that you can view schema objects and data in these databases
- Migrating these databases to Oracle, to take advantage of the full range of Oracle Database features and capabilities

This topic contains the following topics:

- [Migration: Basic Options and Steps](#) (page 2-1)
- [Migration: Background Information and Guidelines](#) (page 2-6)
- [SQL Developer User Interface for Migration](#) (page 2-32)
- [Command-Line Interface for Migration](#) (page 2-34)

2.1 Migration: Basic Options and Steps

To migrate all or part of a third-party database to Oracle, you have the following basic options:

- [Migrating Using the Migration Wizard](#) (page 2-2)
- [Copying Selected Tables to Oracle](#) (page 2-5)

However, before you perform any migration actions, you may want to prepare by setting any appropriate migration user preferences (such as date and timestamp masks and *Is Quoted Identifier On?*) and by reading relevant topics in Migration: Background Information and Guidelines.

After you migrate by using the wizard or by copying tables to Oracle, verify that the results are what you expected.

Tip:

For a walk-through of a typical migration, go to the `sqldeveloper\sqldeveloper\bin` folder and enter the following command:

```
sdcli migration -help=guide
```

Related Topics:

[Migration](#) (page 1-111)

[Migration: Background Information and Guidelines](#) (page 2-6)

2.1.1 Migrating Using the Migration Wizard

The Migration wizard provides convenient, comprehensive guidance through the actions that can be involved in database migration (capturing the source database, converting it to Oracle format, generating DDL to perform the conversion, and so on). This is the recommended approach when performing a migration: you can resolve issues during these phases, and you can then inspect or modify objects to suit your needs.

The migration wizard is invoked in a variety of contexts, such as when you right-click a third-party database connection and select **Migrate to Oracle** or when you click **Tools**, then **Migration**, then **Migrate**. Sometimes the wizard is invoked at a page other than the first step.

On all pages except the last, enabling **Proceed to Summary Page** causes **Next** to go to the Summary page.

2.1.1.1 Repository

The Repository page of the wizard requires that you specify the database connection for the migration repository to be used.

The migration repository is a collection of schema objects that SQL Developer uses to manage metadata for migrations. If you do not already have a migration repository and a database connection to the repository, create them as follows:

1. Create an Oracle user named `MIGRATIONS` with default tablespace `USERS` and temporary tablespace `TEMP`; and grant it at least the `RESOURCE` role and the `CREATE SESSION`, `CREATE VIEW`, and `CREATE MATERIALIZED VIEW` privileges. (For multischema migrations, you must grant the `RESOURCE` role with the `ADMIN` option; and you must also grant this user the `CREATE ROLE`, `CREATE USER`, and `ALTER ANY TRIGGER` privileges, all with the `ADMIN` option.)
2. Create a database connection named `Migration_Repository` that connects to the `MIGRATIONS` user.
3. Right-click the `Migration_Repository` connection, and select **Migration Repository**, then **Associate Migration Repository** to create the repository.

If you do not already have a database connection to the third-party database to be migrated, create one. (For migrations, you should set the third party JDBC driver preference before creating the connection.) For example, create a database connection named `Sales_Sybase` to a Sybase database named `sales`.

Connection: The database connection to the migration repository to be used.

Truncate: If this option is enabled, the repository is cleared (all data from previous migrations is removed) before any data for the current migration is created.

2.1.1.2 Project

The Project page of the wizard specifies the migration project for this migration. A migration project is a container for migration objects.

New lets you create a new project, or **Existing** lets you select from a list of existing projects.

Name: Name to be associated with this migration project.

Description: Optional descriptive comments about the project.

Output Directory: The directory or folder in which all scripts generated by the migration wizard will be placed. Enter a path or click **Choose** to select the location.

2.1.1.3 Source Database

The Source Database page of the wizard specifies the third-party database to be migrated.

Mode: Online causes the migration to be performed by SQL Developer when you have completed the necessary information in the wizard; **Offline** causes SQL Developer to perform the migration using a file (the Offline Capture Source File) that you specify.

Connection (Online mode): The database connection to the third-party database to be migrated. To add a connection to the list, click the Add (+) icon; to edit the selected connection, click the Edit (pencil) icon.

Available Source Platforms (Online mode) List of third-party databases that you can migrate. If the desired platform is not listed, you probably need the appropriate JDBC driver, which you can get by clicking **Help**, then **Check for Updates**, or by clicking the **Add Platform** link and adding the necessary entry on the Database: Third Party JDBC Drivers preferences page.

Offline Capture Source File (Offline mode): The .ocp file. This is a file that you previously created by clicking **Tools**, then **Migration**, then **Create Database Capture Scripts**.

Note:

If you receive the Cannot Connect error, this means that the .ocp file that is normally in the generated offline capture data is not present to identify the type of database, and therefore SQL Developer cannot select the appropriate plugin to perform the conversion. Ensure that the correct, valid .ocp file is present.

2.1.1.4 Capture

The Capture page of the wizard lets you specify the database or databases (of the platform that you specified) to be migrated. Select the desired items under **Available Databases**, and use the arrow icons to move them individually or collectively to **Selected Databases**.

2.1.1.5 Convert

The Convert page of the wizard lets you examine and modify, for each data type in the source database, the Oracle Database data type to which columns of that source type will be converted in the migrated database. For each source data type entry, the possible Oracle Data Type values reflect the valid possible mappings (which might be only one).

Add New Rule: Lets you specify mappings for other source data types.

Edit Rule: Lets you modify the mapping for the selected source data type.

Advanced Options: Displays the Migration: Identifier Options preferences page.

2.1.1.6 Translate

The Translate page of the wizard lets you specify the SQL objects to be translated. Select the desired items under **Available SQL Objects**, and use the arrow icons to move them individually or collectively to **Selected SQL Objects**.

2.1.1.7 Target Database

The Target Database page of the wizard specifies the Oracle database to which the third-party database or databases will be migrated.

Mode: Online causes the migration to be performed by SQL Developer when you have completed the necessary information in the wizard; **Offline** causes SQL Developer to generate scripts after you have completed the necessary information in the wizard, and you must later run those scripts to perform the migration.

Connection: The database connection to the Oracle Database user into whose schema the third-party database or databases are to be migrated. To add a connection to the list, click the Add (+) icon; to edit the selected connection, click the Edit (pencil) icon.

Generated Script Directory: The directory or folder in which migration script files will be generated (derived based on your previous entry for the project Output Directory).

Drop Target Objects: If this option is enabled, any existing database objects in the target schema are deleted before the migration is performed (thus ensuring that the migration will be into an empty schema).

Advanced Options: Displays the Migration: Generation Options preferences page.

2.1.1.8 Move Data

The Move Data page of the wizard lets you specify options for moving table data as part of the migration. Moving the table data is independent of migrating the table definitions (metadata) Note that if you do not want to move the table data, you can specify the mode as Offline and then simply not run the scripts for moving the data.

Mode: Online causes the table data to be moved by SQL Developer when you have completed the necessary information in the wizard; **Offline** causes SQL Developer to generate scripts after you have completed the necessary information in the wizard, and you must later run those scripts if you want to move the data. (Online moves are convenient for moving small data sets; offline moves are useful for moving large volumes of data.)

Connections for online data move: The Source and Target connections for the third-party and Oracle connections, respectively. To add a connection to either list, click the Add (+) icon; to edit the selected connection, click the Edit (pencil) icon.

Truncate Data: If this option is enabled, any existing data in a target (Oracle) table that has the same name as the source table is deleted before the data is moved. If this option is not enabled, any data from a source table with the same name as the corresponding target (Oracle) table is appended to any existing data in the target table.

2.1.1.9 Summary

The Summary page of the wizard provides a summary of your specifications for the project, repository, and actions, in an expandable tree format. If you want to make any changes, go back to the relevant wizard page.

To perform the migration actions that you have specified, click **Finish**.

2.1.2 Copying Selected Tables to Oracle

To copy one or more tables from a third-party database to an Oracle database, you can select the third-party tables and use the **Copy to Oracle** feature. With this approach, you do not need to create or use a migration repository, or to capture and convert objects.

Note that this approach does not perform a complete migration. It only lets you copy the table, and optionally the table data, from the third-party database to an Oracle database. It does *not* migrate or re-create primary and foreign key definitions and most constraints. (Any UNIQUE constraints or default values are not preserved in the copy. NOT NULL constraints are preserved in most cases.) The approach also does not consider any non-table objects, such as procedures.

In addition, this approach supports autoincrement columns only if the INCREMENT BY value is 1, and if the sequence starts at 1 or is adjusted to MAX VAL + 1 at the first call to the trigger.

If these restrictions are acceptable, this approach is fast and convenient. For example, some database owners only need the basic table definitions and the data copied to an Oracle database, after which they can add keys and constraints in the Oracle database using SQL Developer.

To copy selected tables, follow these steps:

1. Create and open a database connection for the third-party database. (For migrations, you should set the third party JDBC driver preference before creating the connection.)
2. In the Connections navigator, expand the display of Tables for the third-party database connection, and select the table or tables to be migrated.

To select multiple tables, use the standard method for individual and range selections (using the Ctrl and Shift keys) as appropriate.

3. Right-click and select **Copy to Oracle**.
4. In the Choose Database for Copy to Oracle dialog box, select the appropriate entries:

Destination Database Name: Database connection to use for copying the selected tables into the Oracle database. (Only Oracle Database connections are shown for selection.)

Include Data: If this option is enabled, any data in the table in the third-party database is copied to the new table after it is created in the Oracle database. If this option is not enabled, the table is created in the Oracle database but no data is copied.

If Table Exists: Specifies what happens if a table with the same name as the one to be copied already exists in the destination Oracle database: **Indicate Error** generates an error and does not perform the copy; **Append** adds the rows from the copied table to the destination Oracle table; **Replace** replaces the data in the destination Oracle table with the rows from the copied table. Note that if the two tables with the same name do *not* have the same column definitions and if Include Data is specified, the data may or may not be copied, depending on whether the source and destination column data types are compatible.

5. To perform the copy operation, click **Apply**.

If a table with the same name as the one to be copied already exists in the destination Oracle database, then:

- If the two tables do not have the same column definitions, the copy is not performed.
- If the two tables have the same column definitions and if Include Data was specified, the data is appended (that is, the rows from the table to be copied are inserted into the existing Oracle table).

2.2 Migration: Background Information and Guidelines

The following topics provide background information and guidelines that are helpful in planning for a database migration:

- [Overview of Migration](#) (page 2-6)
- [Preparing a Migration Plan](#) (page 2-7)
- [Before You Start Migrating: General Information](#) (page 2-12)
- [Before You Start Migrating: Source-Specific Information](#) (page 2-14)
- [Capturing the Source Database](#) (page 2-17)
- [Creating and Customizing the Converted Model](#) (page 2-18)
- [Generating the DDL for the Oracle Schema Objects](#) (page 2-19)
- [Migrating the Data](#) (page 2-19)
- [Making Queries Case Insensitive](#) (page 2-24)
- [Testing the Oracle Database](#) (page 2-25)
- [Deploying the Oracle Database](#) (page 2-29)

Related Topics:

[Migration: Basic Options and Steps](#) (page 2-1)

[SQL Developer User Interface for Migration](#) (page 2-32)

2.2.1 Overview of Migration

An Oracle database provides you with better scalability, reliability, increased performance, and better security than third-party databases. For this reason, organizations migrate from their current database, such as Microsoft SQL Server, Sybase Adaptive Server, or IBM DB2, to an Oracle database. Although database migration can be complicated, SQL Developer enables you to simplify the process of migrating a third-party database to an Oracle database.

SQL Developer captures information from the source database and displays it in the **captured model**, which is a representation of the structure of the source database. This representation is stored in a **migration repository**, which is a collection of schema objects that SQL Developer uses to store migration information.

The information in the repository is used to generate the **converted model**, which is a representation of the structure of the destination database as it will be implemented in the Oracle database. You can then use the information in the captured model and the

converted model to compare database objects, identify conflicts with Oracle reserved words, and manage the migration progress. When you are ready to migrate, you generate the Oracle schema objects, and then migrate the data.

SQL Developer contains logic to extract data from the data dictionary of the source database, create the captured model, and convert the captured model to the converted model.

Using SQL Developer to migrate a third-party database to an Oracle database provides the following benefits:

- Reduces the effort and risks involved in a migration project
- Enables you to migrate an entire third-party database, including triggers and stored procedures
- Enables you to see and compare the captured model and converted model and to customize each if you wish, so that you can control how much automation there is in the migration process

2.2.1.1 Migration Implemented as SQL Developer Extensions

Migration support is implemented in SQL Developer as a set of extensions. If you want, you can disable migration support or support for migrating individual third-party databases.

To view the installed extensions, and to enable or disable individual extensions, click **Tools**, then **Preferences**, then **Extensions**. Note that SQL Developer ships with all extensions and third-party database "plugins" available at the time of release, so to begin migrations, only the third-party drivers need be installed.

Related Topics:

[Manage Features and Updates](#) (page 6-112)

2.2.2 Preparing a Migration Plan

This topic describes the process of how to create a migration plan. It identifies the sections to include in the migration plan, describes how to determine what to include for each section, and explains how to avoid the risks involved in a migration project. This information includes:

- [Task 1: Determining the Requirements of the Migration Project](#) (page 2-7)
- [Task 2: Estimating Workload](#) (page 2-9)
- [Task 3: Analyzing Operational Requirements](#) (page 2-10)
- [Task 4: Analyzing the Application](#) (page 2-11)
- [Task 5: Planning the Migration Project](#) (page 2-11)

2.2.2.1 Task 1: Determining the Requirements of the Migration Project

In this task, you identify which databases you want to migrate and applications that access that database. You also evaluate the business requirements and define testing criteria.

To determine the requirements of the migration project:

1. Define the scope of the project.

There are several choices you must make about the third-party database and applications that access that database in order to define the scope of the migration project. To obtain a list of migration issues and dependencies, you should consider the following

- What third-party databases are you migrating?
 - What is the version of the third-party database?
 - What is the character set of the third-party database?
- What source applications are affected by migrating the third-party database to an Oracle database?
 - What is the third-party application language?
 - What version of the application language are you using?

In the scope of the project, you should have identified the applications you must migrate. Ensure that you have included all the necessary applications that are affected by migrating the database

- What types of connectivity issues are involved in migrating to an Oracle database?
 - Do you use connectivity software to connect the applications to the third-party database? Do you need to modify the connectivity software to connect the applications to the Oracle database?
 - What version of the connectivity software do you use? Can you use this same version to connect to the Oracle database?
- Are you planning to rewrite the applications or modify the applications to work with an Oracle database?

2. Determine whether you have a complex or simple source database environment. Identify the requirements based on the specific scenario.

A **complex scenario** involves more than one of the following:

- Large database (greater than 25 GB)
- Data warehouse
- Large applications (more than 100 forms, reports, and batch jobs)
- Database used by multiple lines of business
- Distributed deployment
- Large user base (more than 100)
- High availability requirement (such as a 24 X 7 X 365 environment)

A **simple scenario** involves the following:

- Small database (less than 25 GB)
- Simple online transaction processing (OLTP)

- Small application (less than 100 forms, reports, and batch jobs)
- Database used by one department
- Centralized deployment
- Small user base (less than 100)
- Average availability (business hours)

If the migration project is a simple scenario, you may not have to complete all possible migration tasks. You make decisions based on your specific environment. For example, if you have a complex scenario, you may require extra testing based on the complexity of the application accessing the database.

3. Determine whether the destination database requires additional hardware and rewriting of backup schedules.
4. Define testing and acceptance criteria.

Define tests to measure the accuracy of the migration. You then use the acceptance criteria to determine whether the migration was successful. The tests that you develop from the requirements should also measure stability, evaluate performance, and test the applications. You must decide how much testing is necessary before you can deploy the Oracle database and applications into a production environment.

5. Create a requirements document with a list of requirements for the migration project.

The requirements document should have clearly defined tasks and number each specific requirement, breaking these into sub-requirements where necessary.

2.2.2.2 Task 2: Estimating Workload

In this task, you use SQL Developer to make calculated decisions on the amount of work that can be automated and how much is manual.

To estimate the workload:

1. Capture the captured model, create the converted model, and migrate to the destination database.

You can analyze the source database through the captured model and a preview of the destination database through the converted model. After you have captured the source database, analyze the captured data contained in the captured model and the converted model. Ensure the content and structure of the migration repository is correct and determine how much time the entire process takes.

2. Use the Migration Log pane to evaluate the capture and migration process, categorize the total number of database objects, and identify the number of objects that can be converted and migrated automatically.

The migration log provides information about the actions that have occurred and record any warnings and errors. They identify the changes that have been made to the converted model so that you can evaluate if you should make changes to the applications that access the destination database.

3. Evaluate and categorize the issues that occurred. The migration log can help by providing information about:

- Tables that did not load when you captured the source database
 - Stored procedures, views, and triggers that did not parse when you created the converted model
 - Syntax that requires manual intervention
 - Database objects that were not created successfully when you migrated the destination database
 - Data that did not migrate successfully when you migrated the destination database
4. For each error or warning in the migration log, evaluate the following:
- Number of times an issue occurred
 - Time required to fix the issues, in person-hours
 - Number of resources required to fix the issue
- After you have solved a complex problem, it should be easier and quicker to resolve the next time you have the same problem.

2.2.2.3 Task 3: Analyzing Operational Requirements

In this task, you analyze the operational requirements, as follows:

1. Evaluate the operational considerations in migrating the source database to a destination database. Consider the following questions:

Note:

If the scope of the migration project is a complex scenario, Oracle recommends that you answer all of these questions. If you have a simple scenario, determine the answers to the most appropriate questions.

- What backup and recovery changes do you require?
 - What downtime is required during the migration?
 - Have you met the performance requirements?
 - Are you changing the operational time window?
 - What effect does the downtime have on the business?
 - What training requirements or additional staff considerations are required?
 - Is it necessary to have the third-party and the Oracle database running simultaneously?
2. For each task, determine the resources and time required to complete.
 3. Create an initial project plan.

Use the information that you have gathered during the requirements and planning stage to develop an initial project plan.

2.2.2.4 Task 4: Analyzing the Application

In this task, you identify the users of the applications that run on the source database, what hardware it requires, what the application does, and how it interfaces with the source database. You also analyze the method the application uses to connect to the database and identify necessary modifications.

Note:

If the migration project is a complex scenario, Oracle recommends that you consider all of the following items. If you have a simple scenario, consider the most relevant items.

To analyze the application:

1. Determine whether changes to the application are required to make them run effectively on the destination database.
2. If changes are required to the application, determine whether it is more efficient to rewrite or modify the applications.

If you are rewriting the application to use the Oracle database, consider the following:

- a. Create the necessary project documentation to rewrite the application. For example, you need a design specification and requirements documentation.
- b. Rewrite the application according to the specification.
- c. Test the application works against the Oracle database.

If you are modifying the application to use the Oracle database, consider the following:

- a. Identify the number of connections to the database that are in the application and modify these connections to use the Oracle database.

You may need to change the connection information to use an ODBC or JDBC connection.

- b. Identify the embedded SQL statements that you need to change in the application before you can test it against the Oracle database.
 - c. Test the application using the Oracle database.
3. Allocate time and resource to address each issue associated with rewriting or modifying the application.
 4. Update the general requirements document for the project that you created in Task 1.

2.2.2.5 Task 5: Planning the Migration Project

In this task, you evaluate the unknown variables that the migration project may contain, such as the difference in the technologies of the source database and the destination database. During the planning stage, you:

- Estimate the budget constraints of the project

- Gather information to produce a migration plan
- Estimate how much time the migration project should take
- Calculate how many resources are required to complete and test the migration

To plan a migration project:

1. Define a list of tasks required to successfully complete the migration project requirements of Task 1.
2. Categorize the list of tasks required to complete the migration project.

You should group these tasks according to your business. This allows you to schedule and assign resources more accurately.
3. Update and finalize the migration project plan based on the information that you have obtained from Task 3 and Task 4.
4. Make sure the migration project plan meets the requirements of the migration project.

The migration plan should include a project description, resources allocated, training requirements, migration deliverable, general requirements, environment analysis, risk analysis, application evaluation, and project schedule.

2.2.3 Before You Start Migrating: General Information

You may need to perform certain tasks before you start migrating a third-party database to an Oracle database. See the following for more information:

- [Creating a Database User for the Migration Repository](#) (page 2-12)
- [Requirements for Creating the Destination Oracle Objects](#) (page 2-13)

See also any information specific to the source database that you will be migrating, as explained in [Before You Start Migrating: Source-Specific Information](#) (page 2-14).

Note:

SQL Developer does not migrate grant information from the source database. The Oracle DBA must adjust (as appropriate) user, login, and grant specifications after the migration.

Note:

Oracle recommends that you make a complete backup of the source database before starting the migration. For more information about backing up the source database, see the documentation for that type of database.

If possible, begin the migration using a development or test environment, not a production database.

2.2.3.1 Creating a Database User for the Migration Repository

SQL Developer requires a migration repository to migrate a third-party database to an Oracle database. To use an Oracle database for the migration repository, you must

have access to that database using a database user account. Oracle recommends that you use a specific user account for migrations. For example, you may want to create a user named MIGRATIONS, create a database connection to that user, and use that connection for the migration repository; and if you wish, you can later delete the MIGRATIONS user to remove all traces of the migration from the database.

When you create a user for migrations, specify the tablespace information as in the following example, instead of using the defaults for tablespaces:

```
CREATE USER migrations IDENTIFIED BY <password>
  DEFAULT TABLESPACE users TEMPORARY TABLESPACE temp,
```

Do not use a standard account (for example, SYSTEM) for migration.

When SQL Developer creates a migration repository, it creates many schema objects that are intended only for its own use. For example, it creates tables, views, indexes, packages, and triggers, many with names starting with MD_ and MIGR. You should not directly modify these objects or any data stored in them.

2.2.3.2 Requirements for Creating the Destination Oracle Objects

The user associated with the Oracle database connection used to perform the migration (that is, to run the script containing the generated DDL statements) must have the following roles and privileges:

Note:

You must grant these privileges directly to a user account. Granting the privileges to a role, which is subsequently granted to a user account, does not suffice. You cannot migrate a database as the user SYS.

Roles

```
CONNECT WITH ADMIN OPTION
RESOURCE WITH ADMIN OPTION
```

Privileges

```
ALTER ANY ROLE
ALTER ANY SEQUENCE
ALTER ANY TABLE
ALTER TABLESPACE
ALTER ANY TRIGGER
COMMENT ANY TABLE
CREATE ANY SEQUENCE
CREATE ANY TABLE
CREATE ANY TRIGGER
CREATE VIEW WITH ADMIN OPTION
CREATE MATERIALIZED VIEW WITH ADMIN OPTION
CREATE PUBLIC SYNONYM WITH ADMIN OPTION
CREATE ROLE
CREATE USER
DROP ANY SEQUENCE
DROP ANY TABLE
DROP ANY TRIGGER
DROP USER
DROP ANY ROLE
GRANT ANY ROLE
INSERT ANY TABLE
```

```
SELECT ANY TABLE  
UPDATE ANY TABLE
```

For example, you can create a user called `migrations` with the minimum required privileges required to migrate a database by using the following commands:

```
CREATE USER migrations IDENTIFIED BY password  
  DEFAULT TABLESPACE users TEMPORARY TABLESPACE temp;  
  
GRANT CONNECT, RESOURCE, CREATE VIEW, CREATE MATERIALIZED VIEW,  
  CREATE PUBLIC SYNONYM TO migrations WITH ADMIN OPTION;  
  
GRANT ALTER ANY ROLE, ALTER ANY SEQUENCE, ALTER ANY TABLE, ALTER TABLESPACE,  
  ALTER ANY TRIGGER, COMMENT ANY TABLE, CREATE ANY SEQUENCE, CREATE ANY TABLE,  
  CREATE ANY TRIGGER, CREATE ROLE, CREATE TABLESPACE, CREATE USER, DROP ANY  
  SEQUENCE, DROP ANY TABLE, DROP ANY TRIGGER, DROP TABLESPACE, DROP USER, DROP ANY  
  ROLE, GRANT ANY ROLE, INSERT ANY TABLE, SELECT ANY TABLE, UPDATE ANY TABLE TO  
  migrations;
```

After you have created the converted model and done first DDL generation done for the new database, it will be clear from the scripts which privileges will be required for your situation.

2.2.4 Before You Start Migrating: Source-Specific Information

Depending on the third-party database that you are migrating to an Oracle database, you may have to configure connection information and install drivers. For more information about specific third-party database requirements, see the following:

- [Before Migrating From IBM DB2](#) (page 2-14)
- [Before Migrating From Microsoft SQL Server or Sybase Adaptive Server](#) (page 2-15)
- [Before Migrating From MySQL](#) (page 2-16)
- [Before Migrating From Teradata](#) (page 2-16)

2.2.4.1 Before Migrating From IBM DB2

To configure an IBM DB2 database for migration:

1. Ensure that the source database is accessible by the IBM DB2 database user that is used by SQL Developer for the source connection. This user must be able to see any objects to be captured in the IBM DB2 database; objects that the user cannot see are not captured. For example, if the user can execute a stored procedure but does not have sufficient privileges to see the source code, the stored procedure cannot be captured.
2. Ensure that you can connect to the IBM DB2 database from the system where you have installed SQL Developer.
3. Ensure that you have downloaded the `db2jcc.jar` and `db2jcc_license_cu.jar` files from IBM.
4. In SQL Developer, do the following:
 - a. Click **Tools**, then **Preferences**, then **Database**, then **Third Party JDBC Drivers**.

- b. Click **Add Entry**.
- c. Select the `db2jcc.jar` file.
- d. Click **OK**.
- e. Repeat steps b through d for the `db2jcc_license_cu.jar` file.

2.2.4.2 Before Migrating From Microsoft SQL Server or Sybase Adaptive Server

To configure a Microsoft SQL Server or Sybase Adaptive Server database for migration:

1. Ensure that the source database is accessible by the Microsoft SQL Server or Sybase Adaptive Server user that is used by SQL Developer for the source connection. This user must be able to see any objects to be captured in the Microsoft SQL Server or Sybase Adaptive Server database; objects that the user cannot see are not captured. For example, if the user can execute a stored procedure but does not have sufficient privileges to see the source code, the stored procedure cannot be captured.
2. Ensure that you can connect to the Microsoft SQL Server or Sybase Adaptive Server database from the system where you have installed SQL Developer.
3. Ensure that you have downloaded the JTDS JDBC driver from <http://sourceforge.net/projects/jtds/>.
4. In SQL Developer, if you have not already installed the JTDS driver using Check for Updates (on the Help menu), do the following:
 - a. Click **Tools**, then **Preferences**, then **Database**, then **Third Party JDBC Drivers**.
 - b. Click **Add Entry**.
 - c. Select the jar file for the JTDS driver you downloaded from <http://sourceforge.net/projects/jtds/>.
 - d. Click **OK**.
5. In SQL Developer, click **Tools**, then **Preferences**, then **Migration: Identifier Options**, and ensure that the setting is correct for the **Is Quoted Identifier On** option (that is, that the setting reflects the database to be migrated).

If this option is enabled, quotation marks (double-quotes) can be used to refer to identifiers; if this option is not enabled, quotation marks identify string literals. As an example of the difference in behavior, consider the following T-SQL code:

```
select col1, "col 2" "column_alias"
from tablex "table_alias"
```

If the **Is Quoted Identifier On** option is enabled (checked), the following PL/SQL code is generated:

```
SELECT col1, col_2 "column_alias"
FROM tablex "table_alias";
```

If the **Is Quoted Identifier On** option is disabled (not checked), the following PL/SQL code is generated:

```
SELECT coll, 'col 2' "column_alias"  
FROM tablex "table_alias";
```

2.2.4.3 Before Migrating From MySQL

To configure a MySQL database for migration, install MySQLConnector/J release 3.1.12 or 5.0.4 on the system where you have installed SQL Developer and set the appropriate SQL Developer preference. Follow these steps:

1. Ensure that you can connect to the MySQL database from the system where you have installed SQL Developer.
2. Ensure that you have downloaded the MySQLConnector/J API from the MySQL website at <http://www.mysql.com/>.
3. In SQL Developer, if you have not already installed the MySQL JDBC driver using Check for Updates (on the Help menu), do the following:
 - a. Click **Tools**, then **Preferences**, then **Database**, then **Third Party JDBC Drivers**.
 - b. Click **Add Entry**.
 - c. Select the jar file for the MySQL driver you downloaded from <http://www.mysql.com/>.
 - d. Click **OK**.
4. Ensure that the source database is accessible by the MySQL user that is used by SQL Developer for the source connection. This user must be able to see any objects to be captured in the MySQL database; objects that the user cannot see are not captured. For example, if the user can execute a stored procedure but does not have sufficient privileges to see the source code, the stored procedure cannot be captured.

2.2.4.4 Before Migrating From Teradata

Note that for the current release of SQL Developer, the following Teradata objects will *not* be migrated to Oracle: procedures, functions, triggers, views, macros, and BTEQ scripts.

To configure a Teradata database for migration:

1. Ensure that the source database is accessible by the Teradata database user that is used by SQL Developer for the source connection. This user must be able to see any objects to be captured in the Teradata database; objects that the user cannot see are not captured.
2. Ensure that you can connect to the Teradata database from the system where you have installed SQL Developer.
3. Ensure that you have downloaded the `tdgssconfig.jar` and `terajdbc4.jar` files from Teradata.
4. In SQL Developer, do the following:
 - a. Click **Tools**, then **Preferences**, then **Database**, then **Third Party JDBC Drivers**.
 - b. Click **Add Entry**.

- c. Select the `tdgssconfig.jar` file.
- d. Click **OK**.
- e. Repeat steps b through d for the `terajdbc4.jar` file.

2.2.5 Capturing the Source Database

Before migrating a third-party database, you must extract information from the database. This information is a representation of the structure of the source database, and it is called the captured model. The process of extracting the information from the database is called capturing the source database.

The capture can be done online or offline:

- **Online capture** is done in a convenient guided sequence, during the migrating using the Migration Wizard process.
- **Offline capture** involves creating a script that you run later. You can use offline capture with IBM DB2, MySQL, Microsoft SQL Server databases, and Sybase Adaptive Server.

After capturing the source database, you can view the source database information in the captured model in SQL Developer. If necessary, you can modify the captured model and change data type mappings.

Note:

Oracle recommends that you do not change the default data type mappings unless you are an experienced Oracle database administrator.

Related Topics:

[Migrating Using the Migration Wizard](#) (page 2-2)

[Offline Capture](#) (page 2-17)

2.2.5.1 Offline Capture

To perform an offline capture of an IBM DB2, MySQL, Microsoft SQL Server, or Sybase Adaptive Server database, you create a set of offline capture scripts, run these scripts outside SQL Developer to create the script output (a dump of the third party metadata tables), and load the script output (the `.ocp` file containing the converted model) using SQL Developer.

- To create the script file (a Windows `.bat` file or a Linux or UNIX `.sh` file) and related files, click **Tools**, then **Migration**, then **Create Database Capture Scripts**.
When this operation completes, you are notified that several files (`.bat`, `.sql`, `.ocp`) have been created, one of which is the controlling script. You must run the controlling script (outside SQL Developer) to populate the object capture properties (`.ocp`) file with information about the converted model.
- To load the converted model from the object capture properties (`.ocp`) file generated by the offline capture controlling script, click **Tools**, then **Migration**, then **Third Party Database Offline Capture**, then **Load Database Capture Script Output**.

2.2.5.1.1 IBM DB2 Offline Capture Notes

Script files and the db2_x.ocp file are generated in the target folder. The main script is startDump.xxx, which you must execute to produce the schema dump. The script files prompt you for the database name, user name, and password, and they use this information to connect to the local DB2 database. The scripts generate the schema dump for database objects within object-specific folders.

To capture the schema information in offline file format, use a command in the following format (with the db2 executable in the run path):

```
db2 -x +o -r <file name> <schema query>
```

To export the schema data in offline file format, use a command in the following format (with the db2 executable in the run path):

- For DB2 version 9 data export:

```
db2 export to <file name> of DEL modified by lobsinsefiles coldel="#"  
timestampformat="\YYYY/MM/DD HH.mm.ss\" datesiso nochardel <select query>
```

- For DB2 version 8 data export:

```
db2 export to <file name> of DEL modified by coldel="#" timestampformat=  
\YYYY/MM/DD HH.mm.ss\" datesiso nochardel <select query>
```

DB2 version 9 supports LOB data in separate files, which is better for migrating large data sizes. With version 8, to support large LOB data, you must modify the oracle ctl file command and db2 command in unload_script.bat or unload_script.sh.

The table data is exported to files with names in the format <catalog>.<schema>.<table>.dat. The format of file is as follows: data1#<COL_DEL>#data2#<COL_DEL>...<ROW_DEL> where COL_DEL and ROW_DEL come from migration offline preference settings.

Before you execute the DB2 data dump script, you must log in by entering a command in the following format:

```
db2 connect to <catalog> user <user name> using <password>
```

You can then execute the script using the logged connection session.

2.2.6 Creating and Customizing the Converted Model

After you capture a third-party database, the next step is to convert it, creating the converted model. The converted model is a representation of the structure of the destination database. SQL Developer creates the converted model using the information from the captured model.

By default, all procedures, functions, triggers, and views are copied to the converted model during translation and translated to Oracle PL/SQL. However, if translation fails for any of the objects, those objects appear in the converted model but their original SQL code remains unchanged. Objects that remain in their original SQL code will not be used when the generation scripts are created. Therefore, to have any such objects migrated, you must either fix the problem in the original SQL code before generating the script or edit the generated script to replace the original SQL code with valid PL/SQL code.

The conversion of the captured model to a converted model is done as part of [Migrating Using the Migration Wizard](#) (page 2-2). You can specify or accept the defaults for data mappings.

Related Topics:

[Correcting Errors in the Converted Model](#) (page 2-19)

2.2.6.1 Correcting Errors in the Converted Model

If error messages with the prefix `Parse Exception` are listed in the migration log, manual intervention is required to resolve the issues. To complete the converted model:

1. Note the converted model schema object that failed.
2. Select that schema object in the converted model.
3. Copy the schema objects DDL and paste it into the translation scratch editor (displayed by clicking Migration, then Translation Scratch Editor).
4. Inspect the properties on the schema object in the translation scratch editor for possible causes of the error.
5. Modify a property of the schema object in the translation scratch editor.
For example, you might comment out one line of a stored procedure.
6. Translate using the appropriate translator.
7. If the error appears again, repeat steps 2 to 6.
8. If the error cannot be resolved in this way, it is best to modify the object manually in the converted model.

2.2.7 Generating the DDL for the Oracle Schema Objects

To generate the DDL statements to create the Oracle schema objects, you must already have captured the captured model and created the converted model. After you generate the DDL, you can run the DDL statements to cause the objects to be created in the Oracle database. At this point, the database schema is migrated to Oracle.

After you generate and run the DDL statements to migrate the schema objects, you can migrate the data from the original source database.

Related Topics:

[Migrating the Data](#) (page 2-19)

2.2.8 Migrating the Data

The Migration Wizard lets you choose whether to migrate (move) any existing data from the source database to the Oracle database. If you choose to migrate the data:

- If you are performing migration in online mode, you can perform the data migration in online or offline mode. (However, for PostgreSQL migrations, the data migration must be performed in online mode.)
- If you are performing the migration in offline mode, the data migration is included in the generated files.

Online data moves are suitable for small data sets, whereas offline data moves are useful for moving large volumes of data.

2.2.8.1 Transferring the Data Offline

To transfer the data offline, you generate and use scripts to copy data from the source database to the destination database. During this process you must:

- Use SQL Developer to generate the data unload scripts for the source database and corresponding data load scripts for the destination database.
- Run the data unload scripts to create data files from the source database using the appropriate procedure for your source database:
 - [Creating Data Files From Microsoft SQL Server or Sybase Adaptive Server](#) (page 2-20)
 - [Creating Data Files From MySQL](#) (page 2-20)
 - For Teradata, perform the offline data move using BTEQ and SQL*Loader.
- Run the data load scripts using SQL*Loader to populate the destination database with the data from these data files.

Related Topics:

[Populating the Destination Database Using the Data Files](#) (page 2-21)

2.2.8.1.1 Creating Data Files From Microsoft SQL Server or Sybase Adaptive Server

To create data files from a Microsoft SQL Server or Sybase Adaptive Server database:

1. Copy the contents of the directory where SQL Developer generated the data unload scripts onto the computer where the source database is installed.
2. Edit the BCP extract script to include the name of the source database server.
 - On Windows, edit the `MicrosoftSqlServer_data.bat` or `Sybase_data.bat` script to alter the bcp lines to include the appropriate variables.

The following shows a line from a sample `MicrosoftSqlServer_data.bat` script:

```
bcp "AdventureWorks.dbo.AWBuildVersion" out "[AdventureWorks].[dbo].[AWBuildVersion].dat" -q -c -t "<EOFD>" -r "<EORD>" -U<Username> -P<Password> -S<ServerName>
```

3. Run the BCP extract script.
 - On Windows, enter one of the following, as appropriate:

```
prompt> MicrosoftSqlServer_data.bat
prompt> Sybase_data.bat
```

This script creates the data files in the current directory.

4. Copy the data files and scripts, if necessary, to the target Oracle database system, or to a system that has access to the target Oracle database and has SQL*Loader (Oracle Client) installed.

2.2.8.1.2 Creating Data Files From MySQL

To create data files from a MySQL database:

1. Copy the contents of the directory where SQL Developer generated the data unload scripts, if necessary, onto the system where the source database is installed or a system that has access to the source database and has the mysqldump tool installed.
2. Edit the `unload_script` script to include the correct host, user name, password, and destination directory for the data files.
 - On Windows, edit the `MySQL_data.bat` script.
 - On Linux or UNIX, edit the `MySQL_data.sh` script.

The following shows a line from a sample `MySQL_data.bat` script:

```
mysqldump -h localhost -u <USERNAME> -p<PASSWORD> -T <DESTINATION_PATH>
--fields-terminated-by="<EOFD>" --fields-escaped-by=" "
--lines-terminated-by="<EORD>" "CarrierDb" "CarrierPlanTb"
```

Edit this line to include the correct values for *USERNAME*, *PASSWORD*, and *DESTINATION PATH*. Do not include the angle brackets in the edited version of this file.

In this command line, `localhost` indicates a loopback connection, which is required by the `-T` option. (See the `mysqldump` documentation for more information.)

3. Run the script.
 - On Windows, enter:


```
prompt> MySQL_data.bat
```
 - On Linux or UNIX, enter:


```
prompt> chmod 755 MySQL_data.sh
prompt> sh ./MySQL_data.sh
```

This script creates the data files in the current directory.
4. Copy the data files and scripts, if necessary, to the target Oracle database system, or to a system that has access to the target Oracle database and has SQL*Loader (Oracle Client) installed.

2.2.8.1.3 Populating the Destination Database Using the Data Files

To populate the destination database using the data files, you run the data load scripts using SQL*Loader:

1. Navigate to the directory where you created the data unload scripts.
2. Edit the `oracle_ctl.bat` (Windows systems) or `oractl_ctl.sh` (Linux or UNIX systems) file, to provide the appropriate user name and password strings.
3. Run the SQL Load script.
 - On Windows, enter:


```
prompt> oracle_ctl.bat
```
 - On Linux or UNIX, enter:


```
prompt> ./oracle_ctl.sh
```

For Microsoft SQL Server and Sybase migrations, if you are inserting into BLOB fields with SQL*Loader, you will receive the following error:

```
SQL*Loader-309: No SQL string allowed as part of LARGEOBJECT field specification
```

To handle situations indicated by this error, you can use either one of the following options:

- Enable the Generate Stored Procedure for Migrate Blobs Offline SQL Developer preference (see : [Migration](#) (page 1-111): Generation Options).
- Use the following Workaround:

Workaround

The workaround is to load the data (which is in hex format) into an additional CLOB field and then convert the CLOB to a BLOB through a PL/SQL procedure.

The only way to export binary data properly through the Microsoft SQL Server or Sybase Adaptive Server BCP is to export it in a hexadecimal (hex) format; however, to get the hex values into Oracle, save them in a CLOB (holds text) column, and then convert the hex values to binary values and insert them into the BLOB column. The problem here is that the HEXTORAW function in Oracle only converts a maximum of 2000 hex pairs. Consequently, write your own procedure that will convert (piece by piece) your hex data to binary. (In the following steps and examples, modify the START.SQL and FINISH.SQL to reflect your environment.

The following shows the code for two scripts, `start.sql` and `finish.sql`, that implement this workaround. Read the comments in the code, and modify any SQL statements as needed to reflect your environment and your needs.

Note:

After you run `start.sql` and before you run `finish.sql`, run BCP; and before you run BCP, change the relevant line in the `.ctl` file from:

```
<blob_column> CHAR(2000000) "HEXTORAW (:<blob_column>)"
```

to:

```
<blob_column>_CLOB CHAR(2000000)
```

```
-- START.SQL
-- Modify this for your environment.

-- This should be executed in the user schema in Oracle that contains the table.
-- DESCRIPTION:
-- ALTERS THE OFFENDING TABLE SO THAT THE DATA MOVE CAN BE EXECUTED
-- DISABLES TRIGGERS, INDEXES AND SEQUENCES ON THE OFFENDING TABLE

-- 1) Add an extra column to hold the hex string;
alter table <tablename> add (<blob_column>_CLOB CLOB);

-- 2) Allow the BLOB column to accept NULLS
alter table <tablename> MODIFY <blob_column> NULL;

-- 3) Disable triggers and sequences on <tablename>
alter trigger <triggername> disable;
```

```
alter table <tablename> drop primary key cascade;

drop index <indexname>;

-- 4) Allow the table to use the tablespace
alter table <tablename> move lob (<blob_column>) store as (tablespace
lob_tablespace);

alter table <tablename> move lob (<blob_column>_clob) store as (tablespace
lob_tablespace);

COMMIT;

-- END OF FILE

-- FINISH.SQL
-- Modify this for your environment.

-- This should be executed in the table schema in Oracle.
-- DESCRIPTION:
-- MOVES THE DATA FROM CLOB TO BLOB
-- MODIFIES THE TABLE BACK TO ITS ORIGINAL SPEC (without a clob)
-- THEN ENABLES THE SEQUENCES, TRIGGERS AND INDEXES AGAIN

-- Currently we have the hex values saved as
-- text in the <blob_column>_CLOB column
-- And we have NULL in all rows for the <blob_column> column.
-- We have to get BLOB locators for each row in the BLOB column

-- put empty blobs in the blob column
UPDATE <tablename> SET <blob_column>=EMPTY_BLOB();

COMMIT;

-- create the following procedure in your table schema
CREATE OR REPLACE PROCEDURE CLOBTOBLOB
AS
inputLength NUMBER; -- size of input CLOB
offSet NUMBER := 1;
pieceMaxSize NUMBER := 2000; -- the max size of each peice
piece VARCHAR2(2000); -- these pieces will make up the entire CLOB
currentPlace NUMBER := 1; -- this is where were up to in the CLOB
blobLoc BLOB; -- blob locator in the table
clobLoc CLOB; -- clob locator pointsthis is the value from the dat file

-- THIS HAS TO BE CHANGED FOR SPECIFIC CUSTOMER TABLE
-- AND COLUMN NAMES
CURSOR cur IS SELECT <blob_column>_clob clob_column , <blob_column> blob_column
FROM /*table*/<tablename> FOR UPDATE;

cur_rec cur%ROWTYPE;

BEGIN

OPEN cur;
FETCH cur INTO cur_rec;

WHILE cur%FOUND
LOOP
--RETRIVE THE clobLoc and blobLoc
```

```
clobLoc := cur_rec.clob_column;
blobLoc := cur_rec.blob_column;

currentPlace := 1; -- reset evertime
-- find the lenght of the clob
inputLength := DBMS_LOB.getLength(clobLoc);

-- loop through each peice
LOOP
-- get the next piece and add it to the clob
piece := DBMS_LOB.subStr(clobLoc,pieceMaxSize,currentPlace);

-- append this piece to the BLOB
DBMS_LOB.WRITEAPPEND(blobLoc, LENGTH(piece)/2, HEXTORAW(piece));

currentPlace := currentPlace + pieceMaxSize ;

EXIT WHEN inputLength < currentplace;
END LOOP;

FETCH cur INTO cur_rec;
END LOOP;

END CLOBtoBLOB;
/

-- now run the procedure
-- It will update the blob column with the correct binary representation
-- of the clob column
EXEC CLOBtoBLOB;

-- drop the extra clob cloumn
alter table <tablename> drop column <blob_column>_clob;

-- 2) apply the constraint we removed during the data load
alter table <tablename> MODIFY FILEBINARY NOT NULL;

-- Now re enable the triggers, indexes and primary keys
alter trigger <triggername> enable;

ALTER TABLE <tablename> ADD ( CONSTRAINT <pkname> PRIMARY KEY ( <column> ) ) ;

CREATE INDEX <index_name> ON <tablename>( <column> );

COMMIT;

-- END OF FILE
```

2.2.9 Making Queries Case Insensitive

With several third-party databases, it is common for queries to be case insensitive. For example, in such cases the following queries return the same results:

```
SELECT * FROM orders WHERE sales_rep = 'Oracle';
SELECT * FROM orders WHERE sales_rep = 'oracle';
SELECT * FROM orders WHERE sales_rep = 'OrAcLe';
```

If you want queries to be case insensitive for a user in the Oracle database, you can create an AFTER LOGON ON DATABASE trigger, in which you set, for that database user, the NLS_SORT session parameter to an Oracle sort name with `_CI` (for "case insensitive") appended.

The following example causes queries for user SMITH to use the German sort order and to be case insensitive:

```
CREATE OR REPLACE TRIGGER set_sort_order AFTER LOGON ON DATABASE
DECLARE
  username VARCHAR2(30);
BEGIN
  username:=SYS_CONTEXT('USERENV','SESSION_USER');
  IF username LIKE 'SMITH' then
    execute immediate 'alter session set NLS_COMP=LINGUISTIC';
    execute immediate 'alter session set NLS_SORT=GERMAN_CI';
  END IF;
END;
```

Related Topics

- Case-Insensitive and Accent-Insensitive Sorts in *Oracle Database 2 Day Developer's Guide*.

2.2.10 Testing the Oracle Database

During the testing phase, you test the application and Oracle database to make sure that the:

- Migrated data is complete and accurate
- Applications function in the same way as the source database
- Oracle database produces the same results as the source database
- Applications and Oracle database meet the operational and performance requirements

You may already have a collection of unit tests and system tests from the original application that you can use to test the Oracle database. You should run these tests in the same way that you ran tests against the source database. However, regardless of added features, you should ensure that the application connects to the Oracle database and that the SQL statements it issues produces the correct results.

Note:

The tests that you run against the application vary depending on the scope of the application. Oracle recommends that you thoroughly test each SQL statement that is changed in the application. You should also test the system to make sure that the application functions the same way as in the third-party database.

Related Topics:

[Testing Methodology](#) (page 2-25)

2.2.10.1 Testing Methodology

Many constraints shape the style and amount of testing that you perform on a database. Testing can contain one or all of the following:

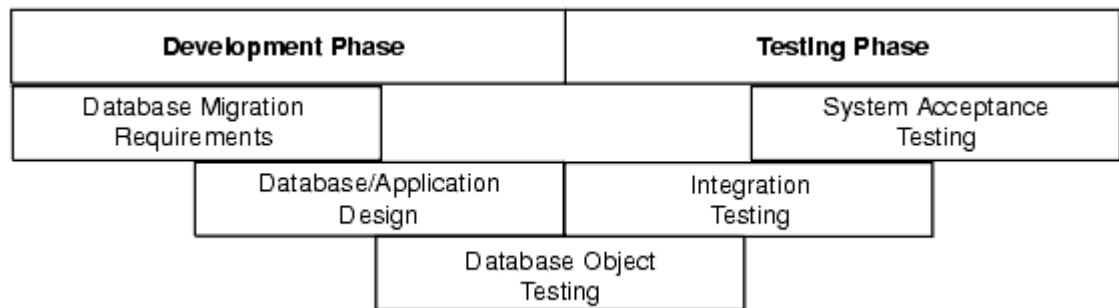
- Simple data validation

- Full life cycle of testing addressing individual unit tests
- System and acceptance testing

You should follow a strategy for testing that suits your organization and circumstances. Your strategy should define the process by which you test the migrated application and Oracle database. A typical test method is the V-model, which is a staged approach where each feature of the database creation is mirrored with a testing phase.

Figure 2-1 (page 2-26) shows an example of the V-model with a database migration scenario:

Figure 2-1 V-model with a Database Migration



There are several types of tests that you use during the migration process. During the testing stage, you go through several cycles of testing to enhance the quality of the database. The test cases you use should make sure that any issues encountered in a previous version of the Oracle database are not introduced again.

For example, if you have to make changes to the migrated schema based on test results, you may need to create a new version of the Oracle database schema. In practice, you use SQL Developer to create a base-line Oracle schema at the start of testing, and then edit this schema as you progress with testing.

Note:

Oracle recommends that you track issues that you find during a testing cycle in an issue tracking system. Track these issues against the version of the database or application that you are testing.

2.2.10.2 Testing the Oracle Database

Use the test cases to verify that the Oracle database provides the same business logic results as the source database.

Note:

Oracle recommends that you define completion criteria so that you can determine the success of the migration.

This procedure explains one way of testing the migrated database. Other methods are available and may be more appropriate to your business requirements.

To test the Oracle database:

1. Create a controlled version of the migrated database.
Oracle recommends that you keep the database migration scripts in a source control system.
2. Design a set of test cases that you can use to test the Oracle database from unit to system level. The test cases should:
 - a. Ensure the following:
 - All the users in the source database have migrated successfully
 - Privileges and grants for users are correct
 - Tables have the correct structure, defaults are functioning correctly, and errors did not occur during mapping or generation
 - b. Validate that the data migrated successfully by doing the following:
 - Comparing the number of rows in the Oracle database with those in the source database
 - Calculating the sum of numerical columns in the Oracle database and compare with those in the source database
 - c. Ensure that the following applies to constraints:
 - You cannot enter duplicate primary keys
 - Foreign keys prevent you from entering inconsistent data
 - Check constraints prevent you from entering invalid data
 - d. Check that indexes and sequences are created successfully.
 - e. Ensure that views migrated successfully by doing the following:
 - Comparing the number of rows in the Oracle database with those in the source database
 - Calculating the sum of numerical columns in the Oracle database and compare with those in the source database
 - f. Ensure that triggers, procedures, and functions are migrated successfully. Check that the correct values are returned for triggers and functions.
3. Run the test cases against the migrated database.
4. Create a report that evaluates the test case results.
These reports allow you to evaluate the data to qualify the errors, file problem reports, and provide a customer with a controlled version of the database.
5. If the tests pass, go to step 7.
If all tests in the test cases pass or contain acceptable errors, the test passes. If acceptable errors occur, document them in an error report that you can use for audit purposes.
6. If the test cases fail:

- a. Identify the cause of the error.
 - b. Identify the test cases needed to check the errors.
 - c. Log an issue on the controlled version of the migrated database code in the problem report.
 - d. Add the test case and a description of the problem to the incident tracking system of your organization, which could be a spreadsheet or bug reporting system. Aside from the test case, the incident log should include the following:
 - Provide a clear, concise description of the incident encountered
 - Provide a complete description of the environment, such as platform and source control version
 - Attach the output of the test, if useful
 - Indicate the frequency and predictability of the incident
 - Provide a sequence of events leading to the incident
 - Describe the effect on the current test, diagnostic steps taken, and results noted
 - Describe the persistent after effect, if any
 - e. Attempt to fix the errors.
 - f. Return to step 1 (page 2-27).
7. Identify acceptance tests that you can use to make sure the Oracle database is an acceptable quality level.

2.2.10.2.1 Guidelines for Creating Tests

You may already have a collection of unit tests and system tests from the original application that you can use to test the Oracle database. However, if you do not have any unit or system tests, you need to create them. When creating test cases, use the following guidelines:

- Plan, specify, and execute the test cases, recording the results of the tests.

The amount of testing you perform is proportional to the time and resources that are available for the migration project. Typically, the testing phase in a migration project can take anywhere from 40% to 60% of the effort for the entire project.
- Identify the components that you are testing, the approach to the test design and the test completion criteria.
- Define each test case so that it is reproducible.

A test that is not reproducible is not acceptable for issue tracking or for an audit process.
- Divide the source database into functions and procedures and create a test case for each function or procedure. In the test case, state what you are going to test, define the testing criteria, and describe the expected results.
- Record the expected result of each test case.

- Verify that the actual results meet the expected results for each test.
- Define test cases that produce negative results as well as those that you expect a positive result.

2.2.10.2.2 Example of a Unit Test Case

The following displays a sample unit test plan for Windows:

Name	Jane Harrison
Module	Table Test Emp
Date test completed	23 May 2007
Coverage log file location	mwb\database\TableTestEmp
Description	This unit test tests that the emp table was migrated successfully.
Reviewed by	John Smith

Task ID	Task Description	Expected Result	Verified (Yes/No)
1	<p>Run the following on the source database for each table:</p> <pre>select count(*) from emp</pre> <p>Run the following on the destination database for each table:</p> <pre>select count(*) from emp</pre>	<p>On the source database, the count(*) produces a number. In this case, the number is the number of rows in each table.</p> <p>On the destination database, the count(*) number corresponds to the number of rows in the new Oracle table.</p>	<p>Yes</p> <p>The number of rows in each table is the same in the source and destination databases.</p>
2	<p>Run the following on the source database for each table:</p> <pre>select sum(salary) from emp</pre> <p>Run the following on the destination database for each table:</p> <pre>select sum(salary) from emp</pre>	<p>On the source database, sum(salary) produces a check sum for the sum of the data in each table.</p> <p>On the destination database, sum(salary) corresponds to the sum of the salary in the emp table.</p>	<p>Yes</p> <p>The sum for each table is the same in the source and destination databases.</p>

2.2.11 Deploying the Oracle Database

Deploying the migrated and tested Oracle database within a business environment can be difficult. Therefore, you may need to consider different rollout strategies depending on your environment. Several rollout strategies are identified for you, but you may use another approach if that is recommended by your organization.

During the deployment phase, you move the destination database from a development to a production environment. A group separate from the migration and testing team, may perform the deployment phase, such as the in-house IT department.

Deployment involves the following:

- [Choosing a Rollout Strategy](#) (page 2-30)

- [Deploying the Destination Database](#) (page 2-30)

2.2.11.1 Choosing a Rollout Strategy

The strategy that you use for migrating a third-party database to an Oracle database must take into consideration the users and the type of business that may be affected during the transition period. For example, you may use the Big Bang approach because you do not have enough systems to run the source database and Oracle database simultaneously. Otherwise, you may want to use the Phased approach to make sure that the system is operating in the user environment correctly before it is released to the general user population. You can use one of the following approaches.

2.2.11.1.1 Phased Approach

Using the Phased approach, you migrate groups of users at different times. You may decide to migrate a department or a subset of the complete user-base. The users that you select should represent a cross-section of the complete user-base. This approach allows you to profile users as you introduce them to the Oracle database. You can reconfigure the system so that only selected users are affected by the migration and unscheduled outages only affect a small percentage of the user population. This approach may affect the work of the users you migrated. However, because the number of users is limited, support services are not overloaded with issues.

The Phased approach allows you to debug scalability issues as the number of migrated users increases. However, using this approach may mean that you must migrate data to and from legacy systems during the migration process. The application architecture must support a phased approach.

2.2.11.1.2 Big Bang Approach

Using the Big Bang approach, you migrate all of the users at the same time. This approach may cause schedule outages during the time you are removing the old system, migrating the data, deploying the Oracle system, and testing that the system is operating correctly. This approach relies on you testing the database on the same scale as the original database. It has the advantage of minimal data conversion and synchronization with the original database because that database is switched off. The disadvantage is that this approach can be labor intensive and disruptive to business activities due to the switch over period needed to install the Oracle database and perform the other migration project tasks.

2.2.11.1.3 Parallel Approach

Using the Parallel approach, you maintain both the source database and destination Oracle database simultaneously. To ensure that the application behaves the same way in the production environment for the source database and destination database, you enter data in both databases and analyze the data results. The advantage of this approach is if problems occur in the destination database, users can continue using the source database. The disadvantage of the Parallel approach is that running and maintaining both the source and the destination database may require more resources and hardware than other approaches.

2.2.11.2 Deploying the Destination Database

There are several ways to deploy the destination database. The following task is an example that you should use as a guideline for deploying the destination database.

Note:

If you have a complex scenario, Oracle recommends that you complete all of the deployment tasks. However, if you have a simple scenario, you should choose the deployment tasks appropriate to your organization.

1. Configure the hardware, if necessary.

In a large scale or complex environment, you must design the disk layout to correspond with the database design. If you use redundant disks, align them in stripes that you can increase as the destination database evolves. You must install and configure the necessary disks, check the memory, and configure the system.

2. Make sure the operating system meets the parameters of the Oracle configuration.

Before installing any Oracle software, make sure that you have modified all system parameters. For more information about modifying system parameters, see the relevant installation guide for your platform, such as Solaris Operating System.

3. Install the Oracle software.

Aside from the Oracle software that allows you to create an Oracle database, you may need to install ancillary software to support the application, such as Extract Transformation and Load (ETL) Software for data warehousing.

4. Create the destination database from the source database and migrate the data to the Oracle database.

There are several ways of putting the destination database into production after testing it, such as:

- Place the successfully tested database into production. The test system is now the production system.
- Use Oracle Export to extract the destination database from the successfully tested database and use Oracle Import to create that database within the production environment.
- Use the tested migration scripts to create the Oracle database and populate it with data using SQL*Loader.

5. Perform the final checks on the destination database and applications.**6. Place the destination database into production using one of the rollout strategies.****7. Perform a final audit by doing the following:**

- Audit the integrity of the data
- Audit the validity of the processes, such as back-up and recovery
- Obtain sign-off for the project, if necessary

Related Topics

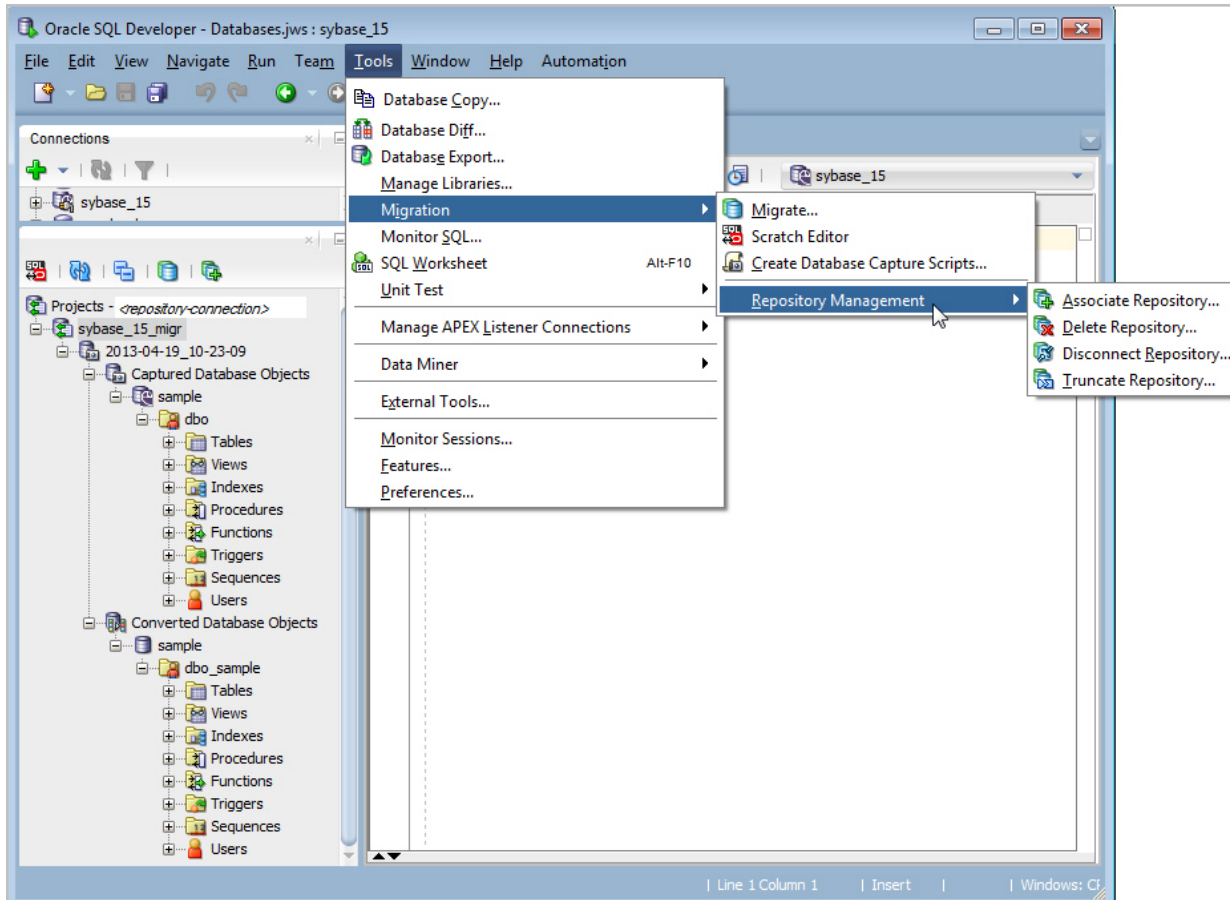
- [Migration: Basic Options and Steps](#) (page 2-1)
- [Migration: Background Information and Guidelines](#) (page 2-6)

- [SQL Developer User Interface for Migration](#) (page 2-32)

2.3 SQL Developer User Interface for Migration

If you are performing database migration, you need to use some migration-specific features. The user interface includes an additional navigator (Migration Projects), a Migration submenu under Tools, and many smaller changes throughout the interface. [Figure 2-2](#) (page 2-32) shows the SQL Developer main window with objects reflecting the migration of a Sybase database. It also shows the Migration submenu.

Figure 2-2 Main Window for a Database Migration



In this figure:

- The Connections navigator shows a connection named `sybase_15`, which is to the Sybase database to be migrated to Oracle. This connection name also appears in a drop-down control in the upper right area.
- In the migration projects navigator, `<repository-connection>` after "Projects -" will be the actual connection name for the migration repository.
- The migration project name is `sybase_15_migr`.
- Under the project name are trees (hierarchies) for Captured Database Objects and Converted Database Objects.

Note:

As an alternative to using the SQL Developer graphical interface for migration tasks, you can use the command line.

2.3.1 Migration Submenu

The Migration submenu contains options related to migrating third-party databases to Oracle. To display the Migration submenu, click **Tools**, then **Migration**.

Migrate: Displays a wizard for performing an efficient migration. The wizard displays steps and options relevant to your specified migration.

Scan Application: Displays the [Application Migration](#) (page 6-2) wizard.

Scratch Editor: Displays the translation scratch editor.

Create Database Capture Scripts specifies options for creating script files, including an offline capture properties (.ocp) file, which you can later load and run.

Repository Management: Enables you to **create (associate)** or **delete** a migration repository, **disconnect** from the current repository (deactivates the current repository but does not disconnect from the database), or **truncate** (remove all data from) the repository.

Related Topics:

[Using the Translation Scratch Editor](#) (page 2-34)

2.3.2 Other Menus: Migration Items

The **View** menu has the following item related to database migration:

- **Migration Projects:** Displays the Migration Projects navigator, which includes any captured models and converted models in the currently selected migration repository.

2.3.3 Migration Preferences

The SQL Developer user preferences window (displayed by clicking **Tools**, then **Preferences**) contains a **Migration** pane with several related subpanes, and a **Translation** pane with a Translation Preferences subpane.

Related Topics:

[Migration](#) (page 1-111)

2.3.4 Migration Log Panes

Migration Log: Contains errors, warnings, and informational messages relating to migration operations.

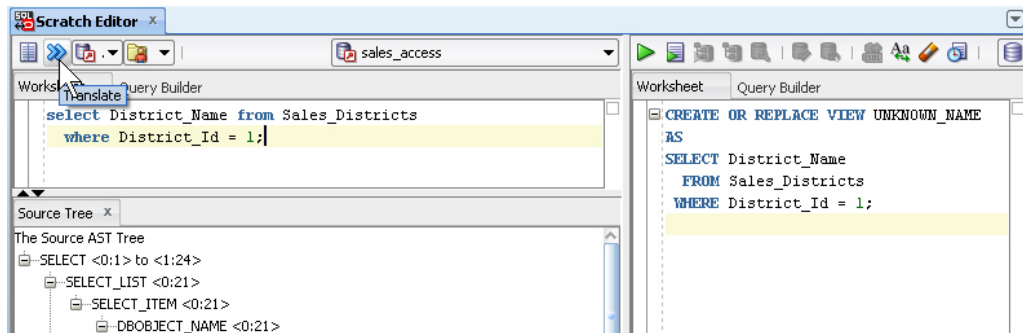
Logging Page: Contains an entry for each migrated-related operation.

Data Editor Log: Contains entries when data is being manipulated by SQL Developer. For example, the output of a Microsoft Excel import operation will be reported here as a series of INSERT statements.

2.3.5 Using the Translation Scratch Editor

You can use the translation scratch editor to enter third-party database SQL statements and have them translated to Oracle PL/SQL statements. You can specify translation from Microsoft SQL Server T-SQL to PL/SQL, or from Sybase T-SQL to PL/SQL.

You can display the scratch editor by clicking **Tools**, then **Migration**, then **Translation Scratch Editor**. The scratch editor consists of two SQL Worksheet windows side by side, as shown in the following figure:



To translate a statement to its Oracle equivalent, select the type of translation, enter the third-party SQL statement or statements; select the specific translation from the Translator drop-down (for example, Access SQL to PL/SQL) and optionally the applicable schema from the Captured Schema drop-down; then click the **Translate** (>>) icon to display the generated PL/SQL statement or statements.

SQL keywords are automatically highlighted.

Note:

For a Microsoft SQL Server or Sybase Adaptive Server connection, the worksheet does not support running T-SQL statements. It only supports SELECT, CREATE, INSERT, UPDATE, DELETE, and DROP statements.

The first time you save the contents of either worksheet window in the translation scratch editor, you are prompted for the file location and name. If you perform any subsequent **Save** operations (regardless of whether you have erased or changed the content of the window), the contents are saved to the same file. To save the contents to a different file, click **File**, then **Save As**.

Related Topics:

[Migration: Basic Options and Steps](#) (page 2-1)

[Migration: Background Information and Guidelines](#) (page 2-6)

[Command-Line Interface for Migration](#) (page 2-34)

2.4 Command-Line Interface for Migration

As an alternative to using the SQL Developer graphical interface for migration operations, you can use the command-line interface, explained in [Command-Line Interface for SQL Developer](#) (page 1-123).

Tip:

For a walk-through of a typical migration, go to the `sqldeveloper\sqldeveloper\bin` folder and enter the following command:

```
sdcli migration -help=guide
```

SQL Developer: Oracle REST Data Services Support

SQL Developer provides a convenient platform for installing and configuring Oracle REST Data Services, and for developing RESTful services.

Topics:

- [Installing Oracle REST Data Services](#) (page 3-1)
- [Oracle REST Data Services Administration](#) (page 3-1)
- [Automatically Enabling REST Access to a Schema, Table, or View \(AutoREST\)](#) (page 3-11)
- [REST Development](#) (page 3-12)

3.1 Installing Oracle REST Data Services

To install Oracle REST Data Services in a traditional (non-CDB) database or a multitenant container database (CDB), click **Tools**, then **REST Data Services**, then **Install**. Follow the instructions for the [Install/Run Oracle REST Data Services](#) (page 6-106) wizard.

To run Oracle REST Data Services in standalone mode, click **Tools**, then **REST Data Services**, then **Run**. Follow the instructions for **Run Standalone Mode** in the [Install/Run Oracle REST Data Services](#) (page 6-106) wizard

To uninstall Oracle REST Data Services from a non-CDB or a CDB, click **Tools**, then **REST Data Services**, then **Uninstall**. Follow the instructions in [Uninstall Oracle REST Data Services](#) (page 6-150).

To reinstall Oracle REST Data Services, you must first uninstall the current installation, and then follow the instructions for the [Install/Run Oracle REST Data Services](#) (page 6-106) wizard.

Related Topics:

[Uninstall Oracle REST Data Services](#) (page 6-150)

3.2 Oracle REST Data Services Administration

The Oracle REST Data Services Administration interface enables you to specify global settings and multiple database settings with different database connections for Oracle REST Data Services. You can create, retrieve, upload, and save REST Data Services settings.

See the following subtopics for more information:

- [About Oracle REST Data Services](#) (page 3-2)
- [Oracle REST Data Services Administration Toolbar and Context Menu](#) (page 3-2)
- [Connecting to Oracle REST Data Services](#) (page 3-3)
- [REST Data Services Navigator: Global Settings and Database Settings](#) (page 3-4)
- [REST Data Services Navigator: NoSQL Store Configuration](#) (page 3-10)
- [REST Data Services Navigator: Reports](#) (page 3-11)

3.2.1 About Oracle REST Data Services

Oracle REST Data Services is a JEE-based alternative for Oracle HTTP Server (OHS) and `mod_plsql`. The JEE implementation offers increased functionality including command line based configuration, enhanced security, file caching and RESTful Web Services. Oracle REST Data Services also provides increased flexibility by supporting deployments using Oracle WebLogic Server, Oracle Glassfish Server, Apache Tomcat, and a standalone mode.

The Oracle Application Express architecture requires some form of Web server to proxy requests between a Web browser and the Oracle Application Express engine. Oracle REST Data Services satisfies this need but its use goes beyond that of Oracle Application Express configurations. Using Oracle REST Data Services simplifies the deployment process because there is no Oracle home required as connectivity is provided using an embedded JDBC driver. See *Oracle REST Data Services Installation, Configuration, and Development Guide* for concepts and details.

3.2.2 Oracle REST Data Services Administration Toolbar and Context Menu

The Oracle REST Data Services Administration window has a toolbar with icons at the top, and a context menu when you right-click the REST Data Services navigator.

Note:

Oracle REST Data Services 3.0 later must be installed and running in order to retrieve or upload administration settings. In addition, the REST Data Services Administrator user must be configured on Oracle REST Data Services. See *Oracle REST Data Services Installation, Configuration, and Development Guide* for information about installing Oracle REST Data Services.

The toolbar and context menu let you perform the following actions:

- **New Administration** displays the Create REST Data Services Administration dialog box, where you specify the **Name** and optionally a **Description** of the administration.

After you click **Apply** in that box, the navigator in the REST Data Services Administration window contains a hierarchical display of global and database settings. Select any node to display the pane for viewing and changing related settings
- **Open File** lets you open a file that contains saved administration settings. If administration settings are already displayed, you will be warned if you want to overwrite the existing settings.

- **Save As** saves the current settings to a .zip file, which you can open later.
- **Retrieve Settings** obtains the administration settings from the REST Data Services.
- **Upload Settings** uploads the administration settings to the REST Data Services. The database settings will be validated before they are uploaded. If any settings are invalid or if required settings are missing, the administration settings will not be uploaded.
Tip: Before you upload the administration settings, click **Test Settings** to validate the database settings. Review the messages and fix any errors.
- **Test Settings** validates the Database Settings and displays informational messages for incorrect settings or required settings that are missing. If errors are not found during validation, a message will be displayed that validation completed successfully.
- **Launch URL** displays the home page in your browser, for example: `http://host:port/apex`
- **Connect** (context menu only): Connects to Oracle REST Data Services.
Retrieve Settings, Upload Settings, and Launch URL are enabled when you connect to Oracle REST Data Services.
- **Disconnect** (context menu only): Disconnects from Oracle REST Data Services.

3.2.3 Connecting to Oracle REST Data Services

Oracle REST Data Services Administration requires a connection to Oracle REST Data Services to retrieve or upload the administration settings. Right-click on the REST Data Services navigator to display the context-menu and select **Connect**. The REST Data Services Connection dialog box is displayed, where you can select a connection, add a connection to the list, or edit a connection in the list.

When you select a connection and click OK, you are always prompted for the REST Data Services Administration password (the password you provided during the installation of REST Data Services). After successful authentication, the administration settings are retrieved and displayed.

If you add or edit a REST Data Services connection, you can specify the following information.

Connection Name: Name to identify the REST Data Services that you are connecting to based on the information that you are providing.

Username: REST Data Services Administrator username that you provided during REST Data Services installation. See *Oracle REST Data Services Installation, Configuration, and Development Guide* for information about configuring the administration user.

http or https: Select which protocol to use.

Hostname: DNS name or IP address of REST Data Services.

Port: Port on which REST Data Services is listening for requests.

Server Path: Location at which REST Data Services is deployed on the application server. Default: /ords (before 2.0.6: /apex)

Note:

If you encounter the **Authentication Failed** error, follow the instructions in [REST Data Services Connection: Authentication Failed](#) (page 6-124).

3.2.4 REST Data Services Navigator: Global Settings and Database Settings

The REST Data Services Administration settings consist of the global settings and database settings for one or more databases. The settings are displayed in the following navigator hierarchy.

```
Administration
  Global Settings
    Connections (page 3-5)
      JDBC
    Default Password (page 3-6)
    Security (page 3-7)
      Allowed Procedures
      Blocked Procedures
      Validation Function
      Virus Scanning
    Cache Files (page 3-9)
      Caching
    Environment (page 3-9)
      Error Reporting
      Log Reporting
      Pre-Processing
      Post-Processing
    Excel (page 3-10)
  Database Settings
    <database-name>
      Connections (page 3-5)
        JDBC
        URL Mapping (Database Settings only)
      ORDS Public User (Database Settings Only) (page 3-7)
      PL/SQL Gateway User (Database Settings Only) (page 3-7)
      APEX REST Users (Database Settings Only) (page 3-7)
```

Global Settings provide the default values for certain Database Settings.

Database Settings specify values in certain categories for REST Data Services connections to specific databases. Oracle REST Data Services supports connecting to multiple databases by letting you create multiple database settings with different database connections and use URL mapping to route requests to the appropriate database.

You can create, edit, or remove multiple database settings in the navigator. To add database settings, right-click Database Settings and select **New Settings**. To remove database settings, right-click on the desired settings and select **Remove Settings**.

In panes that display database settings, any setting values that are highlighted indicate the global settings value. If you change a database settings value and later decide to use the global settings value, you can click **Reset** to restore the global settings values.

To check the validity of database settings, click **Test Settings** in the REST Data Services Administration window toolbar.

3.2.4.1 Connections

Use Connections to specify the connection type: Basic, TNS, or Advanced.

JDBC

Use JDBC to configure JDBC options.

JDBC Driver Type: *thin* or *oci8*.

Initial Pool Size: Initial pool size for the number of connections that will be created.

Maximum Statements: Maximum number of statements to cache for each connection.

Connection Size: Maximum and Minimum number of connections.

Inactivity Timeout: How long (in seconds) an available connection can remain idle before it is closed.

Abandoned Connection Timeout: How long (in seconds) a borrowed (in use) connection can remain unused before it is considered as abandoned and reclaimed.

URL Mappings (Database Settings only)

URL mapping is used in configuring multiple database connections, to define the rules for how requests are routed to the appropriate database.

Specify the rules in the Rule Type and the Routing Rule to your database connection. You can use regular expressions in the Routing Rule.

To add a mapping, click the Add URL (+) icon, select the Rule Type, and specify the Routing Rule. To delete a mapping, select it and click the Remove URL (X) icon.

Rule Type: Type of rule: Request Path, URL Path, or Regular Expression.

Routing Rule: Rule on how to route the request to your database connection. See the examples or routing rules in this topic.

Schema for RESTful Services (optional): The name of the Application Express workspace where RESTful services are defined in the database connection. This can be omitted if RESTful Services are not being used.

Example of routing rules:

URL = `http://example.com/apex/sales/f?p=1:1`

- Protocol: `http`
- Host name: `example.com`
- Context root: `/apex` (location at which the REST Data Services is deployed on the application server)
- Request path: `/sales/f?p=1.1` (the request URL relative to the context root)

Example of routing rule using request path prefix:

Assuming REST Data Services is deployed on `example.com`, its context path is `/apex`, and your database setting name is `sales_prod`, you can create the following request path prefix routing rule:

- Rule Type: Request Path
- Routing Rule: `/sales`

- Schema for RESTful Services: `sales_rest`

The preceding rule means that any requests matching: `http://example.com/apex/sales/...` will be routed to the `sales_prod` database connection. The `sales_rest` schema, specified in the `sales_prod` database connection, will be searched for RESTful Services definitions.

The preceding routing rule *will* match any of the following requests:

```
http://example.com/apex/sales/f?p=1:1
https://example.com/apex/sales/f?p=1:1
http://example.com/apex/sales/leads/
http://www.example.com/apex/sales/forecasting.report?month=jan (if www.example.com
resolves to the same system)
```

The preceding routing rule *will not* match any of the following requests:

```
http://example.com:/apex/leads/f?p=1:1 (missing the /sales prefix)
http://example.com/apex/f?p=1:1 (missing the /sales prefix)
http://example.com/pls/sales/leads/ (the context root is wrong)
```

Example of routing rule using URL prefix:

Assuming REST Data Services is deployed on `example.com`, its context path is `/apex`, and your database setting name is `sales_prod`, you can create the following URL prefix routing rule:

- Rule Type: URL Path
- Routing Rule: `http://example.com/apex/sales`
- Schema for RESTful Services: `sales_rest`

The preceding rule means that any requests matching: `http://example.com/apex/sales/...` will be routed to the `sales_prod` database connection. The `sales_rest` schema, specified in the `sales_prod` database connection, will be searched for RESTful Services definitions.

The preceding routing rule *will* match any of the following requests:

```
http://example.com/apex/sales/f?p=1:1
http://example.com/apex/sales/leads/
http://example.com/apex/sales/forecasting.report?month=jan
```

The preceding routing rule *will not* match any of the following requests:

```
https://example.com/apex/sales/f?p=1:1 (the protocol is wrong)
http://example.com:8080/apex/sales/f?p=1:1 (the port is wrong)
http://example.com/apex/f?p=1:1 (missing the /sales segment of the base URL)
http://example.com/pls/sales/leads/ (the context root is wrong)
http://www.example.com/apex/sales/forecasting.report?month=jan (the host name is
wrong)
```

3.2.4.2 Default Password

You can specify a default password as part of the Global Settings, which means that you can then apply that password in specific Database Settings for any of the following users:

- [ORDS Public User \(Database Settings Only\)](#) (page 3-7)
- [PL/SQL Gateway User \(Database Settings Only\)](#) (page 3-7)

- [APEX REST Users \(Database Settings Only\)](#) (page 3-7)

3.2.4.3 ORDS Public User (Database Settings Only)

Specify a password for the Oracle REST Data Services public user (the database user that has privileges to access the Oracle REST Data Services RESTful services). You cannot change the user name (ORDS_PUBLIC_USER).

To set or reset Password and Confirm Password to the default value specified for default password (in Global Settings), click **Reset**.

3.2.4.4 PL/SQL Gateway User (Database Settings Only)

Specify credentials (user name and password) for the PL/SQL Gateway user, which is responsible for executing stored PL/SQL stored packages and procedures on the specified database connection.

To set or reset Password and Confirm Password to the default value specified for default password (in Global Settings), click **Reset**.

3.2.4.5 APEX REST Users (Database Settings Only)

If you want to be able to use the Oracle Application Express RESTful services, specify the password for the database user APEX_REST_PUBLIC_USER, which used when invoking PL/SQL Gateway operations, and for the database user APEX_LISTENER, which is used to query RESTful services definitions stored in Oracle Application Express.

Specifying these users enables Oracle Application Express applications to publish web services that are accessed by other applications to exchange data. For example, you can configure a RESTful service to return all employee names for a particular department. The data exchange for a RESTful service follows the Representational State Transfer (REST) protocol.

To set or reset Password and Confirm Password to the default value specified for default password (in Global Settings), click **Reset**.

3.2.4.6 Security

Security settings are used to validate procedure requests and determine if a user is allowed to access the procedure.

Verify Secure Socket Layer Requests: Enable this option if HTTPS is available in your environment.

A RESTful service can be protected with the OAuth2 protocol to control access to non-public data. OAuth2 requires all requests involved in the OAuth2 authentication process to be transported via HTTPS. The default behavior of REST Data Services is to verify that all OAuth2 related requests have been received via HTTPS. It will refuse to service any such requests received over HTTP, returning an HTTP status code of 403 Forbidden.

Disabling this option should only be used for development or test environments. Do not disable this option for production environments because it will result in user credentials being passed in clear text.

Security Cache: When you specify a PL/SQL validation function, Oracle REST Data Services uses the security cache to cache the procedures names and determine if they are valid or invalid. Using a caching mechanism reduces the number of database round-trips needed to execute the PL/SQL Validation Function.

Maximum Cache Entries: Maximum cache size. When the maximum entries have been reached, the older (or stale) procedure name that was least recently used is removed from the cache and a new procedure name is added to the cache.

Allowed Procedures

Specify patterns for procedures, packages, or schema names that are allowed to be directly executed from a browser. You can use an asterisk (*) to substitute zero or more characters, and a question mark (?) to substitute for any one character.

A benefit of specifying allowed procedures is that the procedure names and procedure patterns are stored in Oracle REST Data Services. By contrast, invoking the PL/SQL validation function requires a database round-trip.

To add an entry, click the Add Procedure (+) icon and enter the pattern. To delete an entry, select it and click the Remove Procedure (X) icon.

If no allowed procedures are specified, Oracle REST Data Services does not validate the procedure names to determine if the procedure is allowed for processing.

Blocked Procedures

Specify patterns for procedures, packages, or schema names that are forbidden be directly executed from a browser. You can use an asterisk (*) to substitute zero or more characters, and a question mark (?) to substitute for any one character.

If no blocked procedures are specified, Oracle REST Data Services does not validate the procedure names to determine if the procedure is blocked from processing.

To add an entry, click the Add Procedure (+) icon and enter the pattern. To delete an entry, select it and click the Remove Procedure (X) icon.

Disable Default Internal Exclusion List: If option is selected, the Oracle REST Data Services internal exclusion list is not enforced. This internal exclusion list blocks users from accessing the following:

```
sys.*, dbms_*, utl_*, owa_*, owa.*, http.*, htf.*, wpg_docload.*
```

Oracle recommends that you *not* select this option; that is, do not disable the default internal exclusion list. The only possible exception is temporarily disabling the internal exclusion list for debugging purposes.

Validation Function

You can use the Validation Function to determine if the requested procedure in the URL should be allowed for processing. Oracle REST Data Services executes the validation function you specify for each requested procedure, or checks the security cache to determine if a procedure is valid or invalid.

The validation function returns `true` if the procedure is allowed for processing. Otherwise, it returns `false`.

Validation Function Type: Implementation language: PL/SQL or JavaScript.

Validation Function: For PL/SQL, the name of the stored function that resides in the database. If the validation function you specify does not exist, an error message is displayed.

For JavaScript, the code written in JavaScript. The advantage of using JavaScript instead of using the PL/SQL database function is that there are not any database round-trips to execute the validation function.

Virus Scanning

Specify the Internet Content Adaptation Protocol (ICAP) Server name and Port for virus-scanning of your files. A message is displayed if a file is infected.

3.2.4.7 Cache Files

Specify procedure names to allow caching of files.

To add an entry, click the Add Procedure (+) icon and enter the procedure. To delete an entry, select it and click the Remove Procedure (X) icon.

Keep most recently used files: If selected, files that are most recently used will remain in the cache. Specify the maximum number of files to cache. When the maximum entries have been reached, the older files are removed from the cache when a new file is added

Keep files for the specified duration: If selected, files that are cached expire after the specified number of days, hours, or minutes (must be greater than 0).

Caching

Cache File Location: Directory for the cache files.

3.2.4.8 Environment

Application Express Document Table: Name of the document table used by Application Express. (All files uploaded using the PL/SQL Gateway are stored in a document table.)

PL/SQL Gateway

Default Web Page: Default web page to display.

Extension Classes: Extension classes to extend REST Data Services.

Procedure Recorder: If selected, causes procedures to be logged.

Error Reporting

You can specify to show debug messages on the console: or to show error messages in a browser, or to do both or neither. On production systems, you should leave both options unselected (disabled).

Show debug messages on console: If selected, displays debug messages on the console. This may help with problem diagnosis and is appended to the REST Data Services log output. However, you should *not* enable this option on production systems due to the performance impact of outputting large amounts of data to the log.

Show error messages on browser: If selected, displays error responses in the browser, including for each a detailed error message and a stack trace. However, do *not* enable this option on production systems due to the risk of sensitive information being revealed.

Log Reporting

Logging: Specifies whether to make entries in the Oracle REST Data Services log. The log includes activities such as adding a procedure to the cache, finding a procedure in the cache, or reloading a procedure. The log displays the database time and processing time in milliseconds for that procedure.

Keep most recent log messages: If selected, entries that are most recent will remain in the log. Specify the maximum number of log messages to cache. When the maximum entries have been reached, the older log messages are removed from the cache when a new log message is added.

Pre-Processing

Specify the names of procedures to be executed before executing the requested procedure.

To add an entry, click the Add Procedure (+) icon and enter the procedure name. To delete an entry, select it and click the Remove Procedure (X) icon.

Post-Processing

Specify the names of procedures to be executed after executing the requested procedure.

To add an entry, click the Add Procedure (+) icon and enter the procedure name. To delete an entry, select it and click the Remove Procedure (X) icon.

3.2.4.9 Excel

If you are using Oracle Application Express, you have the option of placing your Excel files into an Application Express collection.

Process Excel files into a collection: If selected, puts Excel files into an Application Express Collection.

Put each Excel worksheet into one collection: If selected, puts all Excel worksheets into a single collection, for which you specify the **Collection Name**.

Use the name of the Excel worksheet as the collection name: If selected, creates a collection for each Excel worksheet, and uses each worksheet name for the corresponding collection name.

3.2.5 REST Data Services Navigator: NoSQL Store Configuration

You can create, edit, rename, and delete NoSQL Store configurations that are either secured or non-secured. To create a configuration, right-click NoSQL Store Configuration in the navigator and select Add NoSQL Store.

- To create a configuration, right-click NoSQL Store Configuration in the navigator and select **Add NoSQL Store**.
- To rename a configuration, right-click its name and select **Rename NoSQL Store**.
- To delete a configuration, right-click its name and select **Delete NoSQL Store**, then confirm the deletion.

The display of a NoSQL store configuration includes two tabs: NoSQL Store Configuration and Secure NoSQL Store.

NoSQL Store Configuration tab

Store Alias: The NoSQL store alias, which is used in a NoSQL request that has the following format:

```
host:port/ords/<nosqlstorealias>/<nosqlservice>
```

For more information, see the NoSQL chapter in *Oracle REST Data Services Installation, Configuration, and Development Guide*.

Store Name: Name of the NoSQL store (for example, `kvstore`).

Host and Port: Add and delete host/port combinations.

Role: Add and delete roles.

Maximum number of items per page: The maximum number of items to be returned on each page of the response.

Secure NoSQL Store tab

Secure NoSQL Store: If this option is enabled, the NoSQL store is secure.

Client Security File: Location of the `client.security` file.

Client Trust File: Location of the `client.trust` file.

Username: Name of the user to access the store.

Wallet Directory or Password (and Confirm Password): Method to use for authenticating the specified user.

3.2.6 REST Data Services Navigator: Reports

You can enable reporting that retrieves information from Oracle REST Data Services.

Enable Reporting: If this option is enabled, the Reports page displays information, including statistics, error tracking, and logging.

3.3 Automatically Enabling REST Access to a Schema, Table, or View (AutoREST)

If Oracle REST Data Services has been installed on the system associated with a database connection, and if the connection is open in SQL Developer, you can use the AutoREST feature to conveniently enable or disable Oracle REST Data Services access for specified tables and views in the schema associated with that database connection. Enabling REST access to a table or view allows it to be accessed through RESTful services.

AutoREST is a quick and easy way to expose database tables as REST resources. You sacrifice some flexibility and customizability to gain ease of effort. AutoRest lets you quickly expose data but (metaphorically) keeps you on a set of guide rails. For example, you cannot customize the output formats or the input formats, or do extra validation.

On the other hand, manually created resource modules require you to specify the SQL and PL/SQL to support the REST resources. Using resource modules requires more effort, but offers more flexibility; for example, you can customize what fields are included, do joins across multiple tables, and validate the incoming data using PL/SQL.

So, as an application developer you must make a choice: use the "guide rails" of AutoREST, or create a resource module to do exactly what you need. If you choose AutoREST, you can just enable a table (or set of tables) within a schema.

Note that enabling a schema is not equivalent to enabling all tables and views in the schema. It just means making Oracle REST Data Services aware that the schema exists and that it may have zero or more resources to expose to HTTP. Those resources may be AutoREST resources or resource module resources.

To enable Oracle REST Data Services access to one or more specified tables or views, you must do the following:

1. Enable the schema (the one associated with the connection) for REST access.

Schema level: To enable Oracle REST Data Services access to selected objects (that you specify in the next step) in the schema associated with a connection, right-

click its name in the Connections navigator and select **REST Services**, then **Enable REST Services**.

(To drop support for Oracle REST Data Services access to objects in the schema associated with a connection, right-click its name in the Connections navigator and select **REST Services**, then **Drop REST Services**.)

2. Individually enable REST access for the desired objects.

Table or view level: To enable Oracle REST Data Services access to a specified table or view, right-click its name in the Connections navigator and select **Enable REST Services**.

Follow the instructions in the [RESTful Services Wizard \(Auto-Enable REST Access\)](#) (page 6-126).

3.4 REST Development

You can use SQL Developer to create, maintain, and use RESTful services. See the following topics for more information:

- [About RESTful Services](#) (page 3-12)
- [RESTful Services Terminology](#) (page 3-12)
- [RESTful Services Requirements and Setup](#) (page 3-13)

SQL Developer provides the following options to create and manage RESTful services:

- [REST Development Pane](#) (page 3-14): Use REST Data Services in the Connections navigator to save your RESTful services to the database directly.
- [REST Data Services in the Connections Navigator](#) (page 3-14): Use REST Development using a running Oracle REST Data Services (ORDS), which will save your RESTful services to the database.

3.4.1 About RESTful Services

Representational State Transfer (REST) is a style of software architecture for distributed hypermedia systems such as the World Wide Web. A service is described as RESTful when it conforms to the tenets of REST. A RESTful service has the following characteristics:

- Data is modeled as a set of resources. Resources are identified by URIs
- A small, uniform set of operations are used to manipulate resources (for example, GET, POST, PUT, DELETE).
- A resource can have multiple representations (for example, a blog might have an HTML representation and a RSS representation).
- Services are stateless; and because it is likely that the client will want to access related resources, these should be identified in the representation returned, typically by providing hypertext links.

3.4.2 RESTful Services Terminology

The following are some major terms related to RESTful services:

- **RESTful service:** An HTTP web service that conforms to the tenets of the RESTful architectural style.
- **Resource module:** An organizational unit that is used to group related resource templates.
- **Resource template:** An individual RESTful service that is able to service requests for some set of URIs (Universal Resource Identifiers). The set of URIs is defined by the URI Pattern of the Resource Template
- **URI pattern:** A pattern for the resource template. Can be either a route pattern or a URI template, although you are encouraged to use route patterns.
- **Route pattern:** A pattern that focuses on decomposing the path portion of a URI into its component parts. For example, a pattern of `/objects/:object/:id?` will match `/objects/emp/101` (matches a request for the item in the `emp` resource with `id` of 101) and will also match `/objects/emp/` (matches a request for the `emp` resource, because the `:id` parameter is annotated with the `?` modifier, which indicates that the `id` parameter is optional).

For a detailed explanation of route patterns, see `docs\javadoc\plugin-api\route-patterns.html`, under `<sqldeveloper-install>\ords` and under the location (if any) where you manually installed Oracle REST Data Services.

- **URI template:** A simple grammar that defines the specific patterns of URIs that a given resource template can handle. For example, the pattern `employees/{id}` will match any URI whose path begins with `employees/`, such as `employees/2560`.
- **Resource handler:** Provides the logic required to service a specific HTTP method for a specific resource template. For example, the logic of the GET HTTP method for the preceding resource template might be:


```
select empno, ename, dept from emp where empno = :id
```
- **HTTP operation:** HTTP (HyperText Transport Protocol) defines standard methods that can be performed on resources: GET (retrieve the resource contents), POST (store a new resource), PUT (update an existing resource), and DELETE (remove a resource).

3.4.3 RESTful Services Requirements and Setup

To use the RESTful Services features in SQL Developer, you must install Oracle REST Data Services release 3.0.x.

You can use the Oracle REST Data Services Install wizard to install and run Oracle REST Data Services in standalone mode. See the information for the [Install/Run Oracle REST Data Services](#) (page 6-106) wizard.

You can run the Oracle REST Data Services installer from the command line. For information about using the command-line interface, see *Oracle REST Data Services Installation, Configuration, and Development Guide*.

Note the following about Oracle REST Data Services RESTful services and Oracle Application Express RESTful services:

- To use Oracle REST Data Services RESTful services, you will need to enable a schema for RESTful services. See the [RESTful Services Wizard \(Auto-Enable REST](#)

[Access](#)) (page 6-126) Specify Details page for how to enable a schema for RESTful Services.

- To use Oracle Application Express RESTful services, see *Oracle Application Express Administration Guide* for details to set up Application RESTful services.

3.4.4 REST Data Services in the Connections Navigator

You can create and edit RESTful Services by using the REST Data Services node in the Connections navigator. This is only applicable for Oracle REST Data Services release 3.0.5 and later.

Note:

If you have an Oracle REST Data Services release 3.0.4 or earlier, you can upgrade to a later release.

In the Connections navigator, you can:

- Create, edit, and delete resource modules, resource templates, resource handlers, privileges, and roles saving them directly to the database.
- Export REST definitions for resource modules, privileges, and roles to a SQL file.
- Use the Save REST Handler icon in the SQL Worksheet toolbar to save the REST Handler definitions directly to the database.

Related Topics:

[Installing Oracle REST Data Services](#) (page 3-1)

[RESTful Services Wizard and Object Dialog Boxes](#) (page 6-127)

3.4.5 REST Development Pane

The REST Development pane (click **View**, then **REST Data Services**, then **Development**) enables you to:

- Specify connection information for your RESTful services
When you create a connection, specify the schema alias in the Schema/Workspace field. The schema alias is the Schema Alias name that you provided when you enabled a schema for RESTful Services.
- Retrieve all RESTful services (includes all resource modules and privileges), or retrieve a selected resource module or privilege
- Create, edit, and delete resource modules, resource templates, resource handlers, and privileges
- Upload a resource module or privilege and have the results saved on the server
- Open a .zip file that contain the RESTful services definitions
- Save the RESTful services definitions to a .zip file

The REST Development pane has icons for the following operations:

- **New RESTful Services:** Creates a new RESTful Service that can contain the resource modules and privileges.
- **Retrieve RESTful Services:** Downloads the resource modules and privileges from the user's schema or workspace. (Enabled if you are connected to a RESTful Services connection.)
- **Validate RESTful Services:** Performs validation of all the resource modules' templates and handlers; displays any warning and error messages. (Enabled if the navigator tree contains at least one module or one privilege.)
- **Connect:** Connects to the enabled RESTful Services schema or Application Express workspace based on the user's credentials.

You can right-click on the RESTful Services node to perform the operations with associated icons, plus the following operations:

- **Connect:** Connects to the enabled RESTful Services schema or Application Express workspace based on the user's credentials.
- **Disconnect:** Disconnects from the user's schema or workspace.
- **New RESTful Services:** Creates a new RESTful Service that contains the resource modules and privileges.
- **Retrieve RESTful Services:** Downloads the resource modules and privileges from the user's schema or workspace. (Enabled if you are connected to a RESTful Services connection.)
- **Open File:** Opens a .zip file that contains the resource modules and privileges, and loads them into RESTful Services Administration.
- **Save As:** Creates a .zip file that contains the resource modules and privileges. (Enabled if the navigator tree contains at least one module or one privilege.)

Related Topics:

[RESTful Services Connection](#) (page 6-130)

SQL Developer: Unit Testing

The SQL Developer **unit testing** feature provides a framework for testing PL/SQL objects, such as functions and procedures, and monitoring the results of such objects over time. You create tests, and for each you provide information about what is to be tested and what result is expected. The SQL Developer implementation of unit testing is modeled on the classic and well known xUnit collection of unit test frameworks.

The unit testing feature is part of the support within the SQL Developer family of products for major parts of the life cycle of database system development, from design (provided by Data Modeler) to development to testing.

This topic contains the following topics:

- [Overview of Unit Testing](#) (page 4-1)
- [SQL Developer User Interface for Unit Testing](#) (page 4-3)
- [Unit Test Repository](#) (page 4-4)
- [Editing and Running a Unit Test](#) (page 4-5)
- [Using a Dynamic Value Query to Create a Unit Test](#) (page 4-6)
- [Using Lookups to Simplify Unit Test Creation](#) (page 4-8)
- [Using Variable Substitution in Validation Actions](#) (page 4-10)
- [Unit Test Library](#) (page 4-11)
- [Unit Test Reports](#) (page 4-12)
- [Exporting and Importing Unit Test Objects](#) (page 4-12)
- [Command-Line Interface for Unit Testing](#) (page 4-13)
- [Unit Testing Best Practices](#) (page 4-14)
- [Example of Unit Testing \(Tutorial\)](#) (page 4-16)

4.1 Overview of Unit Testing

The SQL Developer unit testing framework involves a set of sequential steps for each test case. The steps are as follows, including the user input for before the step is run and the framework activities for the step while the test is being run.

1. Identify the object to be tested.

User Input: Identify the object, such as a specific PL/SQL procedure or function.

Framework Activities: Select the object for processing.

2. Perform any startup processing.

User Input: Enter the PL/SQL block, or enter NULL for no startup processing.

Framework Activities: Execute the block.

3. Run the unit test object.

User Input: (None.)

Framework Activities: Execute the unit test.

4. Check and record the results.

User Input: Identify the expected return (result), plus any validation rules.

Framework Activities: Check the results, including for any validation, and store the results.

5. Perform any end processing (teardown).

User Input: Enter the PL/SQL block, or enter NULL for no teardown activities.

Framework Activities: Execute the block.

For each test, you enter the information called for in the preceding steps, to create a **test case**. A **unit test** is a group of test cases (one or more) on a specific PL/SQL object.

Each test case is an **implementation**. Each unit test has at least one implementation (named *Default* by default); however, you can add one or more other implementations. For example, you can have implementations that test various combinations of parameter values, including those that generate exceptions.

When a unit test is run, each implementation is run one after the other. Each implementation runs the startup action (if any) for the test, then the test implementation itself, and then the teardown action (if any). The difference between implementations is in the values of the calling arguments. Any dynamic value query is evaluated before the execution of all of the implementations, including before any startup action.

You can group unit tests into a **test suite** to be run as a grouped item, and the test suite can have its own startup and end processing in addition to any specified for test cases and unit tests.

To learn more about unit testing with SQL Developer, take whichever approach suits your preference:

- Go to [Example of Unit Testing \(Tutorial\)](#) (page 4-16) and follow the steps, and then return to read the remaining conceptual information under SQL Developer: Unit Testing.
- Or
- Read the remaining conceptual information under SQL Developer: Unit Testing, finishing with [Example of Unit Testing \(Tutorial\)](#) (page 4-16).

Related Topics:

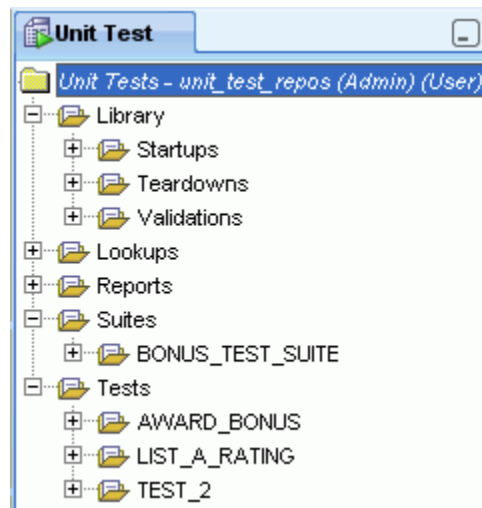
[Using a Dynamic Value Query to Create a Unit Test](#) (page 4-6)

4.2 SQL Developer User Interface for Unit Testing

The SQL Developer user interface for unit testing includes the Unit Test navigator, the Unit Test submenu, and other features.

[Figure 4-1](#) (page 4-3) shows the Unit Test navigator, which includes the top-level nodes Library, Lookups, Reports, Suites, and Tests. (If this navigator is not visible, click **View**, then **Unit Test**.)

Figure 4-1 Unit Test Navigator



In the preceding figure, the top-level repository node shows the name of the connection being used (`unit_test_repos`) and whether the user associated with that connection has only User access to the repository or both Administrator and User access (here, both).

The preceding figure also shows the types of actions under the Library node (Startups, Teardowns, Validations), one test suite, and several tests.

Related Topics:

[Unit Test Submenu](#) (page 4-3)

[Other Menus: Unit Test Items](#) (page 4-4)

[SQL Developer: Unit Testing](#) (page 4-1)

4.2.1 Unit Test Submenu

To display the Unit Test submenu, click **Tools**, then **Unit Test**. (The commands on the Unit Test submenu affect the unit test repository.)

Select Current Repository: Enables you to select the database connection to use for the unit testing repository, and to create a repository using that connection if no repository exists in the associated schema.

Deselect Current Repository: Disconnects from the current unit testing repository. To connect again to a unit testing repository (the same one or a different one), use Select Current Repository.

Purge Run Results: Deletes any existing results from the running of tests and suites.

Create/Update Repository: Enables you to create a unit testing repository, to hold schema objects associated with the SQL Developer unit testing feature.

Drop Repository: Drops (deletes) the current unit testing repository.

Purge Repository Objects: Deletes the contents of the current unit testing repository, but does not delete the repository metadata.

Import From File:

Manage Users: Enables you to select, add, and modify database connections to be used for the unit testing repository.

Show Shared Repository:

Select As Shared Repository: Makes the current repository a shared repository.

Deselect As Shared Repository: Makes the current repository an unshared repository.

Related Topics:

[Unit Test Repository](#) (page 4-4)

4.2.2 Other Menus: Unit Test Items

The **View** menu has the following item related to unit testing:

- **Unit Test:** Toggles the display of the Unit Test navigator.

4.2.3 Unit Test Preferences

The SQL Developer user preferences window (displayed by clicking **Tools**, then **Preferences**) contains a **Unit Test Parameters** pane.

Related Topics:

[Unit Test Parameters](#) (page 1-116)

4.3 Unit Test Repository

The unit test **repository** is a set of tables, views, indexes, and other schema objects that SQL Developer maintains to manage the use of the unit testing feature. (Most of these objects have `UT_` in their names.) You can create a separate database user for a repository or use the schema of an existing database user; but for simplicity and convenience in an environment with a single main shared repository, you may want to create a separate database user.

A repository can be unshared or shared, depending on how many and which database users are allowed to perform various types of unit testing operations:

- In an **unshared repository**, only the database user that owns the unit test repository schema objects can be used for operations that can modify the repository.
There can be multiple unshared repositories, for example, to allow individual developers to create private repositories.
- In a **shared repository**, the owner of the repository objects and any other user that has been granted Administrator access to the repository (specifically, `UT_REPO_ADMINISTRATOR` role) can perform administrative operations, such as managing users.

There can be at most one shared repository, and this is the typical case for a team development environment. A repository administrator can add users and can switch the repository status between shared and unshared. (When a repository is made shared, SQL Developer creates public synonyms for the appropriate repository objects.)

To change an unshared repository to shared, click **Tools**, then **Unit Test**, then **Repository**, then **Select As Shared Repository**. To change a shared repository to unshared, click **Tools**, then **Unit Test**, then **Repository**, then **Deselect As Shared Repository**.

Related Topics:

[Managing Repository Users and Administrators](#) (page 4-5)

[SQL Developer: Unit Testing](#) (page 4-1)

4.3.1 Managing Repository Users and Administrators

To create and run unit tests and suites, you must use a connection for a database user that has been granted User access to the repository (specifically, UT_REPO_USER role). To perform repository administrative operations, such as managing users, you must use a connection for a database user that has been granted Administrator access to the repository (specifically, UT_REPO_ADMINISTRATOR role).

For example, you may want to allow users SCOTT, JONES, and SMITH to use the unit test capabilities and thus have User access to the shared repository, but to allow only SYS and the user that owns the repository objects to have Administrator access to the shared repository.

To grant either type of access to any database users, click **Tools**, then **Unit Test**, then **Repository**, then **Manage Users**. Select the database connection for the owner of the repository objects or for any other user that has been granted Administrator access to the repository. The Unit Testing: Manage Users dialog box is displayed.


Related Topics:

[Unit Testing: Manage Users](#) (page 6-155)

4.4 Editing and Running a Unit Test

To edit or run a unit test, click the unit test name in the Unit Test navigator and select the desired connection for running the unit test. A pane is displayed with two tabs: Details for the unit test specification, and Results for results if you run or debug the test.

The toolbar under the Details tab for the subprogram name has a toolbar that includes the icons shown in the following figure.



- **Freeze View** (the pin) keeps that pane in the SQL Developer window when you click another unit test in the Unit Test navigator; a separate tab and detail view pane are created for that other unit test. If you click the pin again, the unit test's detail view pane is available for reuse.

- **Refresh** refreshes the display in the pane.
- **Debug** starts execution of the first or next implementation of the unit test in debug mode, and displays the results in the Results tab.
- **Run** starts normal execution of the unit test, and displays the results in the Results tab. (Before you click Run, you can specify the database user for the run operation by selecting a database connection on the right.)
- **Edit** (pencil icon) enables you to edit the unit test specification. (If you cannot modify the unit test, lick the Edit icon.)
- **Commit Changes** saves any changes that you have made to the unit test.
- **Rollback Changes** discards any unsaved changes that you have made to the unit test.

If you click the Edit icon, you can modify the Startup Process, Teardown Process, and details for each implementation.

You can also specify **Gather Code Coverage Statistics** to have SQL Developer collect statistics related to code coverage. To view any statistics that have been gathered from unit test runs, use the Test Runs Code Coverage report. In that report, you can click a row with summary information to display detailed information in the Code Coverage Details pane.

Related Topics:

[SQL Developer: Unit Testing](#) (page 4-1)

4.5 Using a Dynamic Value Query to Create a Unit Test

As an alternative to specifying exact input data when creating a unit test, you can create a dynamic value query to use data from a table as input for the test. The query returns values from specified columns in one or more rows, and all sets of values returned are checked by any process validation that you have specified for the test. One common use of dynamic value queries is to perform "reasonableness" tests, such as checking that each salary or price resulting from a test is within a specified range.

To create a test that uses dynamic value queries, create and populate the table to be used by the query, create the test by specifying the object to be tested and any startup and teardown actions, and specify a validation action (such as a query returning rows or no rows).

Note:

A dynamic value query is executed before the execution of all implementations in a test, including any startup action for the test. If you must populate a table before a dynamic value query is evaluated, you can do this is the startup action for a suite that includes the test.

The following example assumes that you have done at least the following in [Example of Unit Testing \(Tutorial\)](#) (page 4-16): created the EMPLOYEES table, created the AWARD_BONUS procedure, and created the unit test repository. It creates a unit test that checks to be sure that no salesperson would receive a bonus so large that his or her salary amount would be greater than 20000. Follow these steps:

1. Create and populate the table for the data by executing the following statements:

```
CREATE TABLE award_bonus_dyn_query (emp_id NUMBER PRIMARY KEY, sales_amt NUMBER);
INSERT INTO award_bonus_dyn_query VALUES (1001, 5000);
INSERT INTO award_bonus_dyn_query VALUES (1002, 6000);
INSERT INTO award_bonus_dyn_query VALUES (1003, 2000);
commit;
```

2. In the Unit Test navigator, right-click the **Tests** node and select **Create Test**.

The Unit Testing: Create Unit Test wizard is displayed. Click **Next** to go from each step to the next; and when you are finished specifying the unit test, click **Finish**.

3. In Select Operation, select the database connection for the schema that you used to create the AWARD_BONUS procedure; then expand the Procedures node and select AWARD_BONUS.
4. In Specify Test Name, for **Test Name** specify AWARD_BONUS_DYN_QUERY (same as the name of the table that you created), and select **Create with single dummy representation**.

5. In Specify Startup, select **Table or Row Copy** because you want to save the current data values in the EMPLOYEES table before any data is modified by the unit test.

When prompted, for **Source Table** specify EMPLOYEES, and for **Target Table** accept the default name provided for a temporary table that will be automatically created when it is needed and deleted when it is no longer needed.

6. In Specify Parameters, click **Next** to go to the next page. (For this example, do not specify the Dynamic Value Query here; instead, you will specify it in later steps.)
7. In Specify Validations, click **Next** to go to the next page.
8. In Specify Teardown, select **Table or Row Restore** because you want to restore the original data values in the EMPLOYEES table before any data was modified by the unit test. When prompted, accept the supplied values for **Target Table** (EMPLOYEES) and **Source Table** (the name of the temporary table).
9. In Summary, review the information. If you need to change anything, click **Back** as needed and make the changes, then proceed to this Summary page. When you are ready to complete the unit test definition, click **Finish**.

10. In the Unit Test navigator, click the node for AWARD_BONUS_DYN_QUERY under Tests, to display the test in an editing window.

11. In the Details pane, click the pencil icon next to Dynamic Value Query, enter the following, and click OK:

```
SELECT emp_id, sales_amt FROM award_bonus_dyn_query;
```

12. For **Expected Result**, leave the value as `Success`.

13. In Specify Validations, click the plus (+) icon and select **Query returning no rows**.

For the query, replace the SELECT statement in the Process Validation box with the following (any semicolon at the end of the statement is ignored):

```
SELECT * FROM employees WHERE salary_amt > 20000
AND commission_pct IS NOT NULL
```

That is, for all salespersons (employees whose commission percentage is not null), check whether the salary resulting from the unit test run is greater than 20000. If there are no such salespersons (that is, if the query returns no rows), the result of the validation action is success.

14. Run the AWARD_BONUS_DYN_QUERY unit test.

Related Topics:

[Unit Testing: Create Unit Test](#) (page 6-152)

[SQL Developer: Unit Testing](#) (page 4-1)

4.6 Using Lookups to Simplify Unit Test Creation

A lookup is an object that contains, for one or more data types, data values that can be tested. Lookups are mainly used for the following purposes:

- Providing lists of values (dropdown lists) for Input fields.
- Automatically creating test implementations based on lookup values.

To create a lookup:

1. In the Unit Test navigator, right-click the **Lookups** node and select **Add Category**.
2. Specify the category name (for example, EMP_ID_LOOKUP).
3. For each data type for which you want to specify lookup values (that is, valid and invalid data values for possible testing), right-click the category name and select **Add Datatype**, select the data type, and use the + (plus sign) icon to add as many data values as you want.

Note that (`null`) is automatically included in the list of values for each data type for each lookup that you create.

For example, for the environment described in [Example of Unit Testing \(Tutorial\)](#) (page 4-16), you could create lookups named EMP_ID_LOOKUP and SALES_AMT_LOOKUP. Each would have only one data type: NUMBER. For the NUMBER data for each lookup, use the + (plus sign) icon to add each of the following values on a separate line, and click the Commit Changes icon or press F11 when you are finished entering the set of numbers for each lookup:

- For EMP_ID_LOOKUP: -100, 99, 1001, 1002, 1003, 1004, 2000, 9999
- For SALES_AMT_LOOKUP: -1000, 0, 1000, 2000, 5000, 6000, 10000, 99999

You can delete and rename lookup categories by using the context (right-click) menu in the Unit Test navigator. You can also delete a data type under a lookup category; however, "deleting" in this case removes any currently specified data values for that type for the lookup category, and it makes the type available for selection in the Unit Testing: Add Data Type dialog box.

Related Topics:

[Providing Values for Input Fields](#) (page 4-9)

[Automatically Creating Implementations](#) (page 4-9)

4.6.1 Providing Values for Input Fields

When you are specifying Input parameters for a unit test implementation, you can click the Lookup Category control to select a lookup category. When you then click in a cell under Input, you can click the dropdown arrow to select a value from the specified lookup. (You can also enter a value other than one in the list.)

For example, if you created the EMP_ID_LOOKUP lookup category, and if you select it as the lookup category when specifying parameters, then the values -100, 99, 1001, 1002, 1003, 1004, 2000, 9999, and (null) will be in the dropdown list for the Input cell for the EMP_ID parameter. (For the SALES_AMT parameter, use the SALES_AMT_LOOKUP category.)

Related Topics:

[Using Lookups to Simplify Unit Test Creation](#) (page 4-8)

4.6.2 Automatically Creating Implementations

If you know that you want implementations to test certain values for a data type, you can use a lookup category to generate these implementations automatically instead of creating them all manually. To do this, use either the DEFAULT lookup category or a user-created category, specify the values for the desired data type, then specify that lookup category for the **Configuration set to use for lookups** preference in the Unit Test Parameters preferences.

For example, assume that for NUMBER input parameters, you always want to check for a very high positive number (such as 9999), a very low negative number (such as -9999), 1, -1, and 0 (zero). Follow these steps:

1. In the Unit Test navigator, expand the **Lookups** node.
2. Right-click DEFAULT and select Add Datatype.
3. In the dialog box, specify NUMBER.
4. In the Lookups Editor for the NUMBER type, use the + (plus sign) icon to add each of the following as a separate item (new line).

```
9999
1.0
0
-1.0
-9999
```

5. Click the Commit Changes icon or press F11.
6. Click **Tools**, then **Preferences**, then **Unit Test Parameters**, and ensure that the configuration set to use for lookups is DEFAULT (the lookup category for which you just specified the values for the NUMBER data type).
7. Create the unit test in the usual way: in the Unit Test navigator, right-click the **Tests** node and select **Create Test**.

However, in the Specify Test Name step, select **Seed/Create implementations using lookup values** (that is, *not* "Create with single dummy representation").

For Specify Startup and Specify Teardown, specify any desired action.

You cannot specify anything for Specify Parameters or Specify Validations now. An implementation (with a name in the form `Test Implementation n`) will automatically be created for each possible combination of input parameters of type NUMBER. For any validation actions, you must specify them later by editing each generated implementation.

Related Topics:

[Unit Test Parameters](#) (page 1-116)

[Unit Testing: Create Unit Test](#) (page 6-152)

[SQL Developer: Unit Testing](#) (page 4-1)

4.7 Using Variable Substitution in Validation Actions

You can use variable substitution in validation actions to write dynamic validations that provide a result based on the values of input and output parameters of a procedure or function, or on the return value of a function. You can specify strings in the following format in validation actions:

- For input parameters: `{PARAMETER_NAME}`
For example, if an input parameter is named `EMP_ID`:

```
SELECT ... WHERE employee_id = {EMP_ID} AND ...;
```
- For output parameters: `{PARAMETER_NAME$}`
For example, if an output parameter is named `SALARY`:

```
SELECT ... WHERE {SALARY$} < old_salary;
```
- For the return value: `{RETURNS$}`
For example, if a function returns a numeric value:

```
SELECT ... WHERE {RETURNS$} > 1;
```

What is actually substituted is the string representation of the parameter value (for text substitution), or the underlying data value of the parameter (for bind substitution, using the syntax `:param-name`). The following example shows both styles of substitution (text style and bind style):

```
DECLARE
  l_PARAM1 DATE;
  bad_date EXCEPTION;
BEGIN
  l_PARAM1 := :PARAM1;
  IF '{PARAM1}' <> TO_CHAR(l_PARAM1)
  THEN
    RAISE bad_date;
  END IF;
END;
```

As a simple example of text-style variable substitution:

- If `P1` is a parameter of type NUMBER and has the value `2.1`, the string `{P1}` will be replaced by the string `2.1`.

- If P1 is a parameter of type VARCHAR2 and has the value ABC, the string '{P1}' will be replaced by the string 'ABC'. (Note the single-quotation marks around {P1} in this example.)

You can use variable substitution for all types of validation actions except Compare Tables. For the applicable validation action types, variable substitution is performed as follows:

- For Query Returning Row(s) and Query Returning No Row(s), substitution is performed on the SQL query.
- For Compare Query Results, substitution is performed on both the source and target SQL queries.
- For Boolean Function and User PL/SQL Code, substitution is performed on the PL/SQL block.

4.8 Unit Test Library

The unit testing library enables you to store actions that you can reuse in the definitions of multiple unit tests. These user-defined actions are displayed under the Library node in the Unit Test navigator. You can store the following kinds of actions in the library, in the following categories:

- Dynamic value queries
- Startup actions
- Teardown actions
- Validation actions

Most categories have subcategories. For example, the *Startup Actions* node has subnodes for *Table or Row Copy* and *User PL/SQL Code*. You can add an entry to the library in the following ways:

- Expand the Library hierarchy to display the relevant lowest-level node (such as *User PL/SQL Code* under *Startups*); right-click and select **Add [action-type]**; specify a name for the action; click the name of the newly created action; and complete the specification.
- Use the **Publish to Library** option when specifying the action when you are creating a unit test: enter a name for the action and click Publish. (The action will be added under the appropriate category and subcategory in the Library display in the Unit Test navigator.)

To use an action from the library when you are creating a unit test, select it from the list under **Library** on the appropriate page in the Unit Testing: Create Unit Test wizard or when you are editing a unit test. When you select an action from the library, you have the following options for incorporating it into the process (startup, teardown, or validation):

- **Copy:** Uses a copy of the action, which you can then edit (for example, to modify the WHERE clause in a User PL/SQL Code procedure). If the action is later changed in the library, it is *not* automatically re-copied into the process.
- **Subscribe:** Uses the action as stored in the library. (You cannot edit the action in the process if you use the Subscribe option.) If the action is later changed in the library, the changed version *is* automatically used in the process.

Related Topics:

[Unit Testing: Create Unit Test](#) (page 6-152)

[SQL Developer: Unit Testing](#) (page 4-1)

4.9 Unit Test Reports

Several SQL Developer reports provide information about operations related to unit testing. These reports are listed in the Unit Test navigator under the Reports node. The available reports include:

- All Suite Runs
- All Test Implementation Runs
- All Test Runs
- Suite Runs Code Coverage
- Suite Test Implementation Runs
- Suite Test Runs
- Test Implementation Runs
- Test Runs Code Coverage
- User Test Runs (test runs grouped by user)

Each unit testing report contains a top pane with a summary information row for each item. To see detailed information about any item, click in its row to display the information in one or more detail panes below the summary information. For example, if you click in a summary row in the All Test Runs report, details about that test run are displayed under the Test Run Details and Most Recent Code Coverage tabs.

Some reports prompt you for bind variables, where you can accept the default values to display all relevant items or enter bind variables to restrict the display.

Related Topics:

[Bind Variables for Reports](#) (page 1-84)

[SQL Developer: Unit Testing](#) (page 4-1)

4.10 Exporting and Importing Unit Test Objects

You can export and import unit tests, suites, and objects that are stored in the library (such as startup, validation, and teardown actions).

Exporting an object causes all dependent objects to be included in the resulting XML file. For example, if you export a suite, the resulting XML file includes all tests in that suite, as well as all startup, validation, and teardown actions within each test in that suite.

To export an object, right-click its name in the Unit Test navigator and select **Export to File**; then specify the location and name for the XML file that will include the definitions of the objects.

Importing unit test objects from an XML file causes all objects in the file to be created in the appropriate places in the Unit Test navigator hierarchy. If an object already exists in the repository with the same name as an object of the same type in the XML file, it is replaced (overwritten) by the object definition in the XML file.

To import unit test objects, click **Tools**, then **Unit Test**, then **Import from File**; then specify the XML file to be used for the import operation.

Related Topics:

[SQL Developer: Unit Testing](#) (page 4-1)

4.11 Command-Line Interface for Unit Testing

As an alternative to using the SQL Developer graphical interface for running unit tests and suites, and exporting and importing unit test objects, you can use the command line.

When running a unit test from the command-line interface, you can use the following parameters:

- `-db <connection name>` specifies the database connection associated with the database user to be used for running the unit test.
- `-repo <connection name>` specifies the database connection associated with the unit testing repository to be used for running the unit test.
- `{-log <0, 1, 2, 3>}` specifies the logging level, where:
 - 0 = No logging (the default).
 - 1 = Report the status.
 - 2 = Report the status and error message.
 - 3 = Report the status, error message, and return ID value.
- `{-return <return id>}` specifies the return ID value, which is used as the primary key in the results table, and which will allow automation tools to query the results from the database.

The following example runs a unit test named `AWARD_BONUS` in a Windows environment where SQL Developer is installed under `C:\`. (Note that test and suite names are case sensitive for the command-line interface.) This example uses the repository connection for user `unit_test_repos` and runs the test as user `fred`.

```
> cd c:\sqldeveloper\sqldeveloper\bin
> sdcli unittest -run -test -name AWARD_BONUS -repo unit_test_repos -db fred
```

The following example exports a unit test named `AWARD_BONUS`. It uses the repository connection for user `unit_test_repos` and stores the exported definitions in the file `C:\ut_xml\award_bonus_test.xml`.

```
> sdcli unittest -exp -test -name AWARD_BONUS -repo unit_test_repos -file c:\ut_xml\award_bonus_test.xml
```

The following example imports object definitions from the file `C:\ut_xml\award_bonus_suite.xml`. It uses the repository connection for user `unit_test_repos`.

```
> sdcli unittest -imp -repo unit_test_repos -file c:\ut_xml\award_bonus_suite.xml
```

To check the results of any tests or suites that you run from the command line, you can start SQL Developer and view the All Test Runs and All Suite Runs reports.

Related Topics:

[Command-Line Interface for SQL Developer](#) (page 1-123)

[Unit Test Reports](#) (page 4-12)

[SQL Developer: Unit Testing](#) (page 4-1)

4.12 Unit Testing Best Practices

This topic contains some recommendations and suggestions for using unit testing in SQL Developer:

- [Strategy](#) (page 4-14)
- [Test Suites](#) (page 4-14)
- [Test Naming](#) (page 4-14)
- [Avoiding Test Naming Clashes](#) (page 4-15)
- [Test Implementations](#) (page 4-15)
- [Library](#) (page 4-15)
- [Lookups](#) (page 4-16)
- [Test and Suite Execution](#) (page 4-16)
- [Reports](#) (page 4-16)

4.12.1 Strategy

If you have many packages, analyze the system as a whole to group the packages into functional areas, and create a test suite for each functional area. This process of decomposition can be done recursively, where the ideal situation is come up with small groups of testable objects that have a common set of argument values.

With a test suite hierarchy in place, you can create tests for each object and place the tests into the hierarchy.

4.12.2 Test Suites

Tests should be organized into test suites to facilitate the bulk execution of tests.

Suites can be built from other suites and tests, allowing areas of interest to be grouped together, and even a "super suite" can be created that executes all tests.

4.12.3 Test Naming

Test names are limited to 120 bytes, so test names can be up to 120 characters for a single-byte character set (significantly smaller for multibyte character sets).

Tests are automatically created using the canonical name of the test object (that is, package functions and procedures will be qualified by the package name). For example, given standalone functions or procedures named MY_PROCEDURE_NAME and MY_FUNCTION_NAME and a package named MY_PACKAGE_NAME, the test

names will be `MY_PACKAGE_NAME.MY_PROCEDURE_NAME` and `MY_PACKAGE_NAME.MY_FUNCTION_NAME`.

4.12.4 Avoiding Test Naming Clashes

If you try to create a test with the same name as an existing test (for example, creating a test multiple times on the same object or on an object with the same name in a different schema), then a sequential number is appended to the new test name. This might result, for example, in the following tests:

```
MY_PACKAGE_NAME.MY_PROCEDURE_NAME
MY_PACKAGE_NAME.MY_PROCEDURE_NAME_1
MY_PACKAGE_NAME.MY_PROCEDURE_NAME_2
```

However, you may want to consider these alternatives:

- If you have objects with the same name in different schemas, it is recommended that you prefix (prepend) the test name with either the physical schema name or a logical synonym. For example, in the following full test names, the last part is the same but the names start with different schema names:

```
USER3.MY_PACKAGE_NAME.MY_PROCEDURE_NAME
USER4.MY_PACKAGE_NAME.MY_PROCEDURE_NAME
USER5.MY_PACKAGE_NAME.MY_PROCEDURE_NAME
```

- If there is a valid reason to add a test for the same object more than once, then it may be better to give each test a distinct name it rather than use the default "sequence" approach. For example:

```
MY_PACKAGE_NAME.MY_PROCEDURE_NAME#LATEST
MY_PACKAGE_NAME.MY_PROCEDURE_NAME#COMPATIBLE
```

4.12.5 Test Implementations

When a test is created, a child Implementation of the test is also created. Each implementation forms the configuration for the execution of a test. The first implementation is named Test Implementation 1. You can create additional Implementations to exercise the object with different combinations of values and environment.

It is recommended that you use implementation names that reflect the test strategies. For example, instead of using the names Test Implementation 1, Test Implementation 2, and Test Implementation 3, use names like Upper Boundaries, Lower Boundaries, and Default Values.

4.12.6 Library

The library is a repository of commonly used values (but not data values, for which you should use [Lookups](#) (page 4-16)). If you find you are entering the same values into the unit testing panels (for example, Startup Process), you can place those values in the library and reuse them in multiple places.

You can take this a step further and ensure that all values are stored in the library, whether they are reused or not. This brings more order to the test building process, and means that as the tested logic changes, it is easy to update all tests accordingly.

4.12.7 Lookups

Lookups store data type value domains organized into categories. For example, a SALES category might have a NUMBER data type with domain values of (-1, 0, 1, 1000, 1000000000).

Categories can be created to group data type values at a fine grain (for example, EMPLOYEE or SET_3) or at the most coarse grain (for example, DEFAULT).

A test implementation can be associated with only one lookup category, so you can choose a category to cover the values for all the implementations of a single test, in which case it is recommended that the lookup name echo the corresponding test name.

4.12.8 Test and Suite Execution

You can execute tests and test suites using the SQL Developer graphical and command-line interfaces. It may be more convenient to use the command-line interface to execute suites or a "super suite" (consisting of all tests).

You could create a generator to run against the UT_TEST and UT_SUITE tables found in the repository schema (or public synonyms for a shared repository) to generate the operating system commands necessary to execute tests and suites.

4.12.9 Reports

The Unit Test navigator contains a set of predefined reports to display test execution results.

You can also run a report, right-click on a grid and select Save Grid as Report, and then view the report to see how the unit testing tables have been built into queries.

4.13 Example of Unit Testing (Tutorial)

This section presents a simplified example in which you create a table and a PL/SQL procedure, create unit tests with test cases for valid and invalid input data, run the unit tests, and create and run a unit test suite. It assumes that you have a table of employee data that includes salary information, and that you need to create a procedure to award bonuses to sales representatives, whose pay consists of a base salary plus a commission-based bonus.

The EMPLOYEES table includes the following columns, all of type NUMBER:

- EMPLOYEE_ID: Employee identification (badge) number.
- COMMISSION_PCT: Commission percentage for the employee: a decimal fraction representing the percentage of the amount of sales by the employee, to be used to compute a bonus that will be added to the employee's base salary to determine the total salary. For example, 0.2 or .2 indicates a 20 percent commission, or 0.2 times the amount of sales.

Only employees in the Sales department have numeric COMMISSION_PCT values. Other employees (not "on commission") have null COMMISSION_PCT values.

- SALARY: Salary amount for the employee; includes base salary plus any bonus (which will be calculated by an award_bonus procedure, to be created during this example).

Assume that the following data exists in these columns in the EMPLOYEES table:

EMPLOYEE_ID	COMMISSION_PCT	SALARY
1001	0.2	8400
1002	0.25	6000
1003	0.3	5000
1004	(null)	10000

You create a procedure named AWARD_BONUS, which has two input parameters:

- `emp_id`: The employee ID of an employee.
- `sales_amt`: The amount of sales with which the employee is credited for the period in question.

This amount is calculated using the COMMISSION_PCT value for the specified employee, and the result is added to the SALARY value for that employee.

If the COMMISSION_PCT is null for the employee, no commission or bonus can be calculated, and an exception is raised. This scenario occurs if an attempt is made to add a commission-based bonus to the salary of an employee who is not in the Sales department.

The rest of this example involves the following major steps:

1. [Create the EMPLOYEES Table](#) (page 4-17).
2. [Create the AWARD_BONUS Procedure](#) (page 4-18).
3. [Create the Unit Testing Repository](#) (page 4-18).
4. [Create a Unit Test](#) (page 4-19).
5. [Run the Unit Test](#) (page 4-20).
6. [Create and Run an Exception Unit Test](#) (page 4-21).
7. [Create a Unit Test Suite](#) (page 4-22).
8. [Run the Unit Test Suite](#) (page 4-23).

4.13.1 Create the EMPLOYEES Table

This tutorial uses a table named EMPLOYEES, which must exist before you run any unit tests of the AWARD_BONUS procedure. This table contains some of the columns used in the HR.EMPLOYEES table that is included in the Oracle-supplied sample schemas, but it does not contain all of the columns, and it contains fewer rows and different data.

You can create this EMPLOYEES table in an existing schema and using an existing database connection, or you can create a new schema and connection for the table. To create and populate this table, enter the following statements in a SQL Worksheet or a SQL*Plus command window:

```
-- Connect as the database user that will be used to run the unit tests.
-- Then, enter the following statements:
```

```
CREATE TABLE employees (employee_id NUMBER PRIMARY KEY, commission_pct NUMBER,
salary NUMBER);
INSERT INTO employees VALUES (1001, 0.2, 8400);
INSERT INTO employees VALUES (1002, 0.25, 6000);
INSERT INTO employees VALUES (1003, 0.3, 5000);
-- Next employee is not in the Sales department, thus is not on commission.
INSERT INTO employees VALUES (1004, null, 10000);
commit;
```

4.13.2 Create the AWARD_BONUS Procedure

Create the AWARD_BONUS procedure in the same schema as the EMPLOYEES table. In a SQL Worksheet using the appropriate database connection, enter the following text:

```
create or replace
PROCEDURE award_bonus (
  emp_id NUMBER, sales_amt NUMBER) AS
  commission REAL;
  comm_missing EXCEPTION;
BEGIN
  SELECT commission_pct INTO commission
  FROM employees
  WHERE employee_id = emp_id;

  IF commission IS NULL THEN
    RAISE comm_missing;
  ELSE
    UPDATE employees
    SET salary = salary + sales_amt*commission
    WHERE employee_id = emp_id;
  END IF;
END award_bonus;
/
```

Click the Run Script icon (or press F5) to create the AWARD_BONUS procedure.

4.13.3 Create the Unit Testing Repository

You will need a unit testing repository in the database to hold schema objects that you create and that SQL Developer will maintain. You can create a separate database user for this repository or use the schema of an existing database user; however, to simplify your learning and any possible debugging you may need to do later, it is recommended that you use a separate schema for the unit testing repository, and the instructions in this section reflect this approach.

1. Create a database user (for example, UNIT_TEST_REPOS) for the unit testing repository. Using a database connection with DBA privileges, right-click **Other Users** in the Connections navigator and select **Create User**. Specify UNIT_TEST_REPOS as the user name, and complete any other required information.

For **Default Tablespace**, specify USERS; for **Temporary Tablespace**, specify TEMP.

For **System Privileges**, enable CREATE SESSION; then click **Apply**, then **Close**.

2. Create a database connection for the unit testing repository user that you created, as follows. Click **Tools**, then **Unit Test**, then **Manage Users**. In the Select

Connection dialog box, click the plus (+) icon to create a new database connection (for example, unit_test_repos) for the unit testing repository user.

Click **Save** to save the connection, then **Cancel** to close the dialog box.

3. Create the repository in the schema of the user that you created, as follows. Click **Tools**, then **Unit Test**, then **Select Current Repository**. Specify the database connection (for example, unit_test_repos) for the unit testing repository user. When you see a message that no repository exists for that connection, follow the prompts to create a new repository.

SQL Developer will display several prompts so it can execute commands that grant the necessary privileges to the unit test repository user. In each case, click **Yes**, and enter the SYS account password when prompted.

4.13.4 Create a Unit Test

To create the first unit test, use the Unit Test navigator. If this navigator is not visible on the left side, click **View**, then **Unit Test**.

1. In the Unit Test navigator, right-click the **Tests** node and select **Create Test**.
The Unit Testing: Create Unit Test wizard is displayed. In the remaining steps, click **Next** to go from each step to the next; and when you are finished specifying the unit test, click **Finish**.
2. In **Select Operation**, select the database connection for the schema that you used to create the AWARD_BONUS procedure; then expand the **Procedures** node and select AWARD_BONUS.
3. In **Specify Test Name**, for **Test Name** specify AWARD_BONUS (same as the procedure name), and select **Create with single dummy representation**.
4. In **Specify Startup**, click the plus (+) icon to add a startup action; and for the action select **Table or Row Copy** because you want to save the current data values in the EMPLOYEES table before any data is modified by the unit test.

When prompted, for **Source Table** specify EMPLOYEES, and for **Target Table** accept the default name provided for a temporary table that will be automatically created when it is needed and deleted when it is no longer needed. (The target table will be created; and if a table already exists with the name that you specify as the target table, *it will be overwritten.*)

5. In **Specify Parameters**, change the values in the **Input** column to the following:
For Parameter EMP_ID: 1001
For Parameter SALES_AMT: 5000
For **Expected Result**, leave the value as `Success`.
6. In **Specify Validations**, click the plus (+) icon and select **Query returning row(s)**.
For the query, replace the SELECT statement in the **Process Validation** box with the following (any semicolon at the end of the statement is ignored):

```
SELECT * FROM employees
WHERE employee_id = 1001 AND salary = 9400
```

That is, because employee 1001 has a 20 percent (0.2) commission and because the sales amount was specified as 5000, the bonus is 1000 (5000 * 0.2), and the new

salary for this employee is 9400 (8400 base salary plus 1000 bonus). In this case, the query returns one row, and therefore the result of the validation action is success.

Note that you could have instead specified the SELECT statement in this step using variable replacement, as follows:

```
SELECT * FROM employees
WHERE employee_id = {EMP_ID} AND salary = 9400
```

However, in this specific example scenario, using variable substitution would provide no significant advantage.

7. In Specify Teardown, select **Table or Row Restore** because you want to restore the original data values in the EMPLOYEES table before any data was modified by the unit test. When prompted, accept the supplied values for **Target Table** (EMPLOYEES) and **Source Table** (the name of the temporary table).
8. In Summary, review the information. If you need to change anything, click **Back** as needed and make the changes, then proceed to this Summary page. When you are ready to complete the unit test definition, click **Finish**.

Related Topics:

[Unit Testing: Create Unit Test](#) (page 6-152)

[SQL Developer User Interface for Unit Testing](#) (page 4-3)

4.13.5 Run the Unit Test

To run the unit test, use the Unit Test navigator. If this navigator is not visible on the left side, click **View**, then **Unit Test**.

1. In the Unit Test navigator, expand the **Tests** node and click the AWARD_BONUS test.
A pane for the AWARD_BONUS test is displayed, with Details and Results tabs.
2. On the Details tab, near the top-right corner, select the database connection for the schema that you used to create the AWARD_BONUS procedure.
Do not change any other values. (However, if you later want to run the unit test with different specifications or data values, you can click the Edit (pencil) icon in the Code Editor toolbar at the top of the pane.)
3. Click the Run Test (green arrowhead) icon in the Code Editor toolbar (or press F9).

At this point, focus is shifted to the Results tab, where you can soon see that the AWARD_BONUS ran successfully.

If you want to check the EMPLOYEES table data, you will see that the salary for employee 1001 is the same as it was before (8400), because the startup action for the unit test copied the original data to the temporary table and the teardown action restored the original data to the EMPLOYEES table.

Related Topics:

[SQL Developer User Interface for Unit Testing](#) (page 4-3)

4.13.6 Create and Run an Exception Unit Test

Create another unit test for the exception condition where the COMMISSION_PCT value is null for the employee, and therefore no commission or bonus can be calculated. For this tutorial, the test data includes employee 1004 with a null commission percentage. (This condition could result from several possible scenarios, the most likely being an attempt to run the procedure on a salaried employee who is not eligible for commissions.)

The steps for creating this exception unit test are similar to those in [Create a Unit Test](#) (page 4-19), except there are no startup or teardown steps because this test should not modify any table data, and there is no need for any validation action.

1. In the Unit Test navigator, right-click the **Tests** node and select **Create Test**.

The Unit Testing: Create Unit Test wizard is displayed. Click **Next** to go from each step to the next; and when you are finished specifying the unit test, click **Finish**.

2. In Select Operation, select the database connection for the schema that you used to create the AWARD_BONUS procedure; then expand the Procedures node and select AWARD_BONUS.
3. In Specify Test Name, for **Test Name** specify AWARD_BONUS_NO_COMM_EXC, and select **Create with single dummy representation**.
4. In Specify Startup, click **Next** to go to the next page.
5. In Specify Parameters, change the values in the Input column to the following:
EMP_ID: 1004
SALES_AMT: 5000
For **Expected Result**, change the value to `Exception` and leave the expected error number as `ANY`.
6. In Specify Validations, click **Next** to go to the next page.
7. In Specify Teardown, click **Next** to go to the next page.
8. In Summary, review the information. If you need to change anything, click **Back** as needed and make the changes, then proceed to this Summary page. When you are ready to complete the unit test definition, click **Finish**.

To run this unit test, follow the steps in [Run the Unit Test](#) (page 4-20), except specify AWARD_BONUS_NO_COMM_EXC instead of AWARD_BONUS.

On the Results tab, you will see that the AWARD_BONUS_NO_COMM_EXC test ran successfully; and if you check the EMPLOYEES table data, you will see that the information for employee 1004 (and all the other employees) was not changed.

Note:

As an alternative to creating a separate unit test for the exception condition, you could add it as an **implementation** to the AWARD_BONUS test (right-click AWARD_BONUS and select **Add Implementation**). Thus, the AWARD_BONUS unit test would have two implementations: the "Default" implementation using employee 1001, and the AWARD_BONUS_NO_COMM_EXC implementation using employee 1004.

The approach in this tutorial enables you to create a simple unit test suite using the two unit tests. However, in more realistic unit testing scenarios, it is probably better to use a unit test for each procedure, add implementations for each test case for a procedure, and group multiple unit tests (for individual procedures) into one or more test suites.

Related Topics:

[Create a Unit Test Suite](#) (page 4-22)

[Unit Testing: Create Unit Test](#) (page 6-152)

4.13.7 Create a Unit Test Suite

Create a unit test suite that groups together the two unit tests of the AWARD_BONUS procedure. If the Unit Test navigator is not visible on the left side, click **View**, then **Unit Test**.

1. In the Unit Test navigator, right-click the **Suites** node and select **Add Suite**.
2. In the Unit Testing: Add Test Suite dialog box, specify AWARD_BONUS_SUITE as the suite name.
3. In the Unit Test navigator, under Suites, click the AWARD_BONUS_SUITE node. An pane for the AWARD_BONUS_SUITE test suite is displayed.
4. Do not specify a Startup Process or Teardown Process, because neither is needed for this test suite.
5. Click the Add (+) icon to add the first test to the suite.
6. In the Unit Testing: Add Tests or Suites to a Suite dialog box, click (select) AWARD_BONUS, check (select) **Run Test Startups** and **Run Test Teardowns** so that the startup and teardown actions for that unit test will be run, and click **OK**.
7. Click the Add (+) icon to add the next test to the suite.
8. In the Unit Testing: Add Tests or Suites to a Suite dialog box, click (select) AWARD_BONUS_NO_COMM_EXC, and click **OK**. (The check **Run Test Startups** and **Run Test Teardowns** options are irrelevant here because the AWARD_BONUS_NO_COMM_EXC test does not perform any startup and teardown actions.)
9. Click the Commit Changes icon in the Code Editor toolbar at the top of the pane (or press F11).

Related Topics:

[Unit Testing: Add Tests or Suites to a Suite](#) (page 6-152)

[SQL Developer User Interface for Unit Testing](#) (page 4-3)

4.13.8 Run the Unit Test Suite

To run the unit test suite, use the Unit Test navigator. If you are in the editing pane for the AWARD_BONUS_SUITE test suite, run the suite by clicking the Run Suite (green arrowhead) icon in the Code Editor toolbar. Otherwise, perform the following steps:

1. In the Unit Test navigator, expand the **Suites** node and click the AWARD_BONUS_SUITE test suite.

A pane for the AWARD_BONUS_SUITE test is displayed, with Details and Results tabs.

2. In the Details tab, near the top-right corner, select the database connection for the schema that you used to create the AWARD_BONUS procedure.

Do not change any other values. (However, if you later want to run the unit test suite with different specifications, you can click the Edit (pencil) icon in the Code Editor toolbar at the top of the pane.)

3. Click the Run Suite (green arrowhead) icon in the Code Editor toolbar (or press F9).

After the suite is run, focus is shifted to the Results tab, where you can soon see that the AWARD_BONUS_SUITE test suite ran successfully.

Related Topics:

[Command-Line Interface for Unit Testing](#) (page 4-13)

[SQL Developer: Unit Testing](#) (page 4-1)

SQL Developer Tutorial: Creating Objects for a Small Database

In this tutorial, you will use SQL Developer to create objects for a simplified library database, which will include tables for books, patrons (people who have library cards), and transactions (checking a book out, returning a book, and so on).

Note:

Many other SQL Developer tutorials are available. For information, see the Start Page (click **Help**, then **Start Page**, then the **Get Started** tab if it is not already visible), especially the Tutorials and Online Demonstrations.

The tables are deliberately oversimplified for this tutorial. They would not be adequate for any actual public or organizational library. For example, this library contains only books (not magazines, journals, or other document formats), and it can contain no more than one copy of any book.

You will perform the following major steps:

1. [Create a Table \(BOOKS\)](#) (page 5-2).
2. [Create a Table \(PATRONS\)](#) (page 5-4).
3. [Create a Table \(TRANSACTIONS\)](#) (page 5-5).
4. [Create a Sequence](#) (page 5-7).
5. [Insert Data into the Tables](#) (page 5-7).
6. [Create a View](#) (page 5-8).
7. [Create a PL/SQL Procedure](#) (page 5-9).
8. [Debug a PL/SQL Procedure](#) (page 5-10) (optional).
9. [Use the SQL Worksheet for Queries](#) (page 5-12) (optional).

Note:

To delete the objects that you create for this tutorial, you can use the DROP statements at the beginning of the script in [Script for Creating and Using the Library Tutorial Objects](#) (page 5-13).

This tutorial assumes that you have a database connection to a database user that will own the tables and other objects that you create.

- If you need to create such a database user, do that first (see [Users \(Other Users\)](#) (page 1-26)).
- If you need to create a connection to that user, create the connection (or automatically generate local connections), as explained in [Database Connections](#) (page 1-28).

To print this tutorial in the online help, right-click its book icon in the Help Contents display and select **Print Topic Subtree**.

Related Topics:

[Script for Creating and Using the Library Tutorial Objects](#) (page 5-13)

[SQL Developer Concepts and Usage](#) (page 1-1)

[SQL Developer User Interface](#) (page 1-3)

[Database Objects](#) (page 1-13)

5.1 Create a Table (BOOKS)

The BOOKS table contains a row for each book in the library. It includes columns of character and number types, a primary key, a unique constraint, and a check constraint. You will use the Create Table dialog box to create the table declaratively; the table that you create will be essentially the same as if you had entered the following statement using the SQL Worksheet:

```
CREATE TABLE books (  
    book_id VARCHAR2(20),  
    title VARCHAR2(50)  
        CONSTRAINT title_not_null NOT NULL,  
    author_last_name VARCHAR2(30)  
        CONSTRAINT last_name_not_null NOT NULL,  
    author_first_name VARCHAR2(30),  
    rating NUMBER,  
    CONSTRAINT books_pk PRIMARY KEY (book_id),  
    CONSTRAINT rating_1_to_10 CHECK (rating IS NULL OR  
        (rating >= 1 and rating <= 10)),  
    CONSTRAINT author_title_unique UNIQUE (author_last_name, title));
```

To create the BOOKS table, connect to the database as the user in the schema you want to use for this tutorial. Right-click the Tables node in the schema hierarchy on the left side, select **New Table**, and enter the following information. (If a tab or field is not mentioned, do not enter anything for it. Be sure that the Advanced box is not checked when you start creating the table.)

Schema: Specify your current schema as the schema in which to create the table.

Name: BOOKS

Create the table columns using the following information. After creating each column except the last one (rating), click **Add Column** to add the next column. (If you accidentally click OK instead of Add Column, right-click the BOOKS table in the Connections navigator display, select Edit, and continue to add columns.)

Column Name	Type	Size	Other Information and Notes
book_id	VARCHAR 2	20	Primary Key (Automatically checks Not Null; an index is also created on the primary key column. This is the Dewey code or other book identifier.)
title	VARCHAR 2	50	Not Null
author_last_name	VARCHAR 2	30	Not Null
author_first_name	VARCHAR 2	30	
rating	NUMBER		(Librarian's personal rating of the book, from 1 (poor) to 10 (great))

After you have entered the last column (rating), check **Advanced** (next to Schema). This displays a pane for more table options. For this table, you will use the Unique Constraints and Check Constraints panes.

Unique Constraints pane

Click **Add** to add a unique constraint for the table, namely, that the combination of author_last_name and title must be unique within the table. (This is deliberately oversimplified, since most major libraries will have allow more than one copy of a book in their holdings. Also, the combination of last name and title is not always a "foolproof" check for uniqueness, but it is sufficient for this simple scenario.)

Name: author_title_unique

In **Available Columns**, double-click TITLE and then AUTHOR_LAST_NAME to move them to Selected Columns.

Check Constraints pane

Click **Add** to add a check constraint for the table, namely, that the rating column value is optional (it can be null), but if a value is specified, it must be a number from 1 through 10. You must enter the condition using SQL syntax that is valid in a CHECK clause (but do not include the CHECK keyword or enclosing parentheses for the entire CHECK clause text).

Name: rating_1_to_10

Condition: rating is null or (rating >= 1 and rating <= 10)

Click **OK** to finish creating the table.

Go to [Create a Table \(PATRONS\)](#) (page 5-4) to create the next table.

See Also:

For detailed information about the table dialog box and its tabs, see [Create Table \(quick creation\)](#) (page 6-41) and [Create/Edit Table \(with advanced options\)](#) (page 6-42)

5.2 Create a Table (PATRONS)

The PATRONS table contains a row for each patron who can check books out of the library (that is, each person who has a library card). It includes an object type (MDSYS.SDO_GEOMETRY) column. You will use the Create Table dialog box to create the table declaratively; the table that you create will be essentially the same as if you had entered the following statement using the SQL Worksheet:

```
CREATE TABLE patrons (
  patron_id NUMBER,
  last_name VARCHAR2(30)
    CONSTRAINT patron_last_not_null NOT NULL,
  first_name VARCHAR2(30),
  street_address VARCHAR2(50),
  city_state_zip VARCHAR2(50),
  location MDSYS.SDO_GEOMETRY,
  CONSTRAINT patrons_pk PRIMARY KEY (patron_id));
```

The use of single city_state_zip column for all that information is not good database design; it is done here merely to simplify your work in the tutorial.

The location column (Oracle Spatial and Graph geometry representing the patron's geocoded address) is merely to show the use of a complex (object) type.

To create the PATRONS table, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the Tables node in the schema hierarchy on the left side, select **New Table**, and enter the following information. (If a tab or field is not mentioned, do not enter anything for it. Be sure that the Advanced box is not checked when you start creating the table.)

Schema: Specify your current schema as the schema in which to create the table.

Name: PATRONS

Create most of the table columns using the following information. After creating each column except the city_state_zip column, click **Add Column** to add the next column. (If you accidentally click OK instead of Add Column, right-click the PATRONS table in the Connections navigator display, select Edit, and continue to add columns.)

Column Name	Type	Size	Other Information and Notes
patron_id	NUMBER		Primary Key. (Unique patron ID number, with values to be created using a sequence that you will create)
last_name	VARCHAR2	30	Not Null
first_name	VARCHAR2	30	
street_address	VARCHAR2	30	
city_state_zip	VARCHAR2	30	

The last column in the table (location) requires a complex data type, for which you must use the Columns tab with advanced options. Check **Advanced** (next to Schema). This displays a pane for selecting more table options.

In the Columns pane, click the city_state_zip column name, and click the Add Column (+) icon to add the following as the last column in the table.

Column Name	Type	Other Information and Notes
location	Complex type Schema: MDSYS Type: SDO_GEOMETRY	(Oracle Spatial and Graph geometry object representing the patron's geocoded address)

After you have entered the last column (location), click **OK** to finish creating the table. Go to [Create a Table \(TRANSACTIONS\)](#) (page 5-5) to create the next table.

5.3 Create a Table (TRANSACTIONS)

The TRANSACTIONS table contains a row for each transaction involving a patron and a book (for example, someone checking a book out or returning a book). It includes two foreign key columns. You will use the Create Table dialog box to create the table declaratively; the table that you create will be essentially the same as if you had entered the following statement using the SQL Worksheet:

```
CREATE TABLE transactions (
  transaction_id NUMBER,
  patron_id CONSTRAINT for_key_patron_id
    REFERENCES patrons(patron_id),
  book_id CONSTRAINT for_key_book_id
    REFERENCES books(book_id),
  transaction_date DATE
    CONSTRAINT tran_date_not_null NOT NULL,
  transaction_type NUMBER
    CONSTRAINT tran_type_not_null NOT NULL,
  CONSTRAINT transactions_pk PRIMARY KEY (transaction_id));
```

To create the TRANSACTIONS table, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the Tables node in the schema hierarchy on the left side, select **New Table**, and enter the following information. (If a tab or field is not mentioned, do not enter anything for it. Be sure that the Advanced box is not checked when you start creating the table.)

Schema: Specify your current schema as the schema in which to create the table.

Name: TRANSACTIONS

Create the table columns using the following information. After creating each column except the last one (transaction_type), click **Add Column** to add the next column. (If you accidentally click OK instead of Add Column, right-click the TRANSACTIONS table in the Connections navigator display, select Edit, and continue to add columns.)

Column Name	Type	Size	Other Information and Notes
transaction_id	NUMBER		Primary Key. (Unique transaction ID number, with values to be created using a trigger and sequence that will be created automatically)
patron_id	NUMBER		(Foreign key; must match a patron_id value in the PATRONS table)
book_id	VARCHAR	20 2	(Foreign key; must match a book_id value in the BOOKS table)

Column Name	Type	Size	Other Information and Notes
transaction_date	DATE		(Date and time of the transaction)
transaction_type	NUMBER		(Numeric code indicating the type of transaction, such as 1 for checking out a book)

After you have entered the last column (transaction_type), check **Advanced** (next to Schema). This displays a pane for selecting more table options. For this table, you will use the Column Sequences and Foreign Keys panes.

Column Sequences pane

You have already specified TRANSACTION_ID as the primary key, and you will use this pane only to specify that the primary key column values are to be populated automatically. This convenient approach uses a trigger and a sequence (both created automatically by SQL Developer), and ensures that each transaction ID value is unique.

Column: TRANSACTION_ID

Sequence: New Sequence

Trigger: TRANSACTIONS_TRG (The default; a before-insert trigger with this name will be created automatically.)

Foreign Keys tab

1. Click **Add** to create the first of the two foreign keys for the TRANSACTIONS table.

Name: for_key_patron_id

Referenced Schema: Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers. Use the schema you have been using for this tutorial.

Referenced Table: PATRONS

Referenced Constraint: PATRONS_PK (The name of the primary key constraint for the PATRONS table. Be sure that the **Referenced Column on PATRONS** displayed value is PATRON_ID.)

Associations: Local Column: PATRON_ID

Associations: Referenced Column on PATRONS: PATRON_ID

2. Click **Add** to create the second of the two foreign keys for the TRANSACTIONS table.

Name: for_key_book_id

Referenced Schema: Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers. Use the schema you have been using for this tutorial.

Referenced Table: BOOKS

Referenced Constraint: BOOKS_PK (The name of the primary key constraint for the BOOKS table. Be sure that the **Referenced Column on BOOKS** displayed value is BOOK_ID.)

Associations: Local Column: BOOK_ID

Associations: Referenced Column on BOOKS: BOOK_ID

3. Click **OK** to finish creating the table.

You have finished creating all the tables. To create a sequence for use in generating unique primary key values for the PATRONS table, go to [Create a Sequence](#) (page 5-7).

5.4 Create a Sequence

Create one sequence object, which will be used in INSERT statements to generate unique primary key values in the PATRONS table. (You do not need to create a sequence for the primary key in the TRANSACTIONS table, because you used the SQL Developer feature that enables automatic population of primary key values for that table.) You will use the Create Sequence dialog box to create the sequence declaratively; the sequence that you create will be essentially the same as if you had entered the following statements using the SQL Worksheet:

```
CREATE SEQUENCE patron_id_seq
  START WITH 100
  INCREMENT BY 1;
```

After creating the sequence, you can use it in INSERT statements to generate unique numeric values. The following example uses the patron_id_seq sequence in creating a row for a new patron (library user), assigning her a patron ID that is the next available value of the patron_id_seq sequence:

```
INSERT INTO patrons VALUES (patron_id_seq.nextval,
  'Smith', 'Jane', '123 Main Street', 'Mytown, MA 01234', null);
```

To create the sequence, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the Sequences node in the schema hierarchy on the left side, select **New Sequence**, and enter information using the Create Sequence dialog box.

Schema: Specify your current schema as the schema in which to create the sequence.

Name: patron_id_seq

Increment: 1

Start with: 100

Min value: 100

Click **OK** to finish creating the sequence.

To insert sample data into the tables, go to [Insert Data into the Tables](#) (page 5-7).

5.5 Insert Data into the Tables

For your convenience in using the view and the PL/SQL procedure that you will create, add some sample data to the BOOKS, PATRONS, and TRANSACTIONS tables. (If you do not add sample data, you can still create the remaining objects in this tutorial, but the view and the procedure will not return any results.)

Go to the SQL Worksheet window associated with the database connection you have been using. Copy and paste the following INSERT statements into the **Enter SQL Statement** box:

```
INSERT INTO books VALUES ('A1111', 'Moby Dick', 'Melville', 'Herman', 10);
INSERT INTO books VALUES ('A2222', 'Get Rich Really Fast', 'Scammer', 'Ima', 1);
INSERT INTO books VALUES ('A3333', 'Finding Inner Peace', 'Blissford', 'Serenity',
null);
```

```
INSERT INTO books VALUES ('A4444', 'Great Mystery Stories', 'Whodunit', 'Rodney', 5);
INSERT INTO books VALUES ('A5555', 'Software Wizardry', 'Abugov', 'D.', 10);

INSERT INTO patrons VALUES (patron_id_seq.nextval,
    'Smith', 'Jane', '123 Main Street', 'Mytown, MA 01234', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
    'Chen', 'William', '16 S. Maple Road', 'Mytown, MA 01234', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
    'Fernandez', 'Maria', '502 Harrison Blvd.', 'Sometown, NH 03078', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
    'Murphy', 'Sam', '57 Main Street', 'Mytown, MA 01234', null);

INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (100, 'A1111', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (100, 'A2222', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (101, 'A3333', SYSDATE, 3);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (101, 'A2222', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (102, 'A3333', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (103, 'A4444', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (100, 'A4444', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (102, 'A2222', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (102, 'A5555', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
    transaction_date, transaction_type)
VALUES (101, 'A2222', SYSDATE, 1);
```

Click the Run Script icon, or press the F5 key.

To create a view, go to [Create a View](#) (page 5-8).

Related Topics:

[Using the SQL Worksheet](#) (page 1-41)

5.6 Create a View

Create a view that returns information about patrons and their transactions. This view queries the PATRONS and TRANSACTIONS tables, and returns rows that contain a patron's ID, last name, and first name, along with a transaction and the transaction type. The rows are ordered by patron ID, and by transaction type within patron IDs.

To create the patrons_trans_view view, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the

Views node in the schema hierarchy on the left side, select **New View**, and enter the following information. (If a tab or field is not mentioned, do not enter anything for it.)

Schema: Specify your current schema as the schema in which to create the view.

Name: patrons_trans_view

SQL Query tab

In the SQL Query box, enter (or copy and paste) the following statement:

```
SELECT p.patron_id,
       p.last_name,
       p.first_name,
       t.transaction_type,
       t.transaction_date
FROM patrons p, transactions t
WHERE p.patron_id = t.patron_id
ORDER BY p.patron_id, t.transaction_type
```

Then click **Test Syntax**, and ensure that you have not made any syntax errors. If you made any errors, correct them and click Test Syntax again.

DDL

Review the SQL statement that SQL Developer will use to create the view. If you want to make any changes, go back to the SQL Query tab and make the changes there.

If you want to save the CREATE VIEW statement to a SQL script file, click **Save** and specify the location and file name.

When you are finished, click **OK**.

You have finished creating the view. If you inserted data to the underlying tables, you can see the data returned by this view as follows: in the Connections navigator, expand Views, and select PATRONS_TRANS_VIEW, then click the Data tab.

To create a procedure that lists all books with a specified rating, go to [Create a PL/SQL Procedure](#) (page 5-9).

Related Topics:

[Insert Data into the Tables](#) (page 5-7)

5.7 Create a PL/SQL Procedure

Create a procedure that lists all books with a specified rating. You can then call this procedure with an input parameter (a number from 1 to 10), and the output will be all the titles of all books with that rating.

To create the procedure, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the Procedures node in the schema hierarchy on the left side, select **New Procedure**, and enter the following information using the Create PL/SQL Procedure dialog box.

Object Name: list_a_rating

Click **OK**. A source window for the new procedure is opened. Enter (or copy and paste) the following procedure text, replacing any existing text:

```
CREATE OR REPLACE
PROCEDURE list_a_rating(in_rating IN NUMBER) AS
    matching_title VARCHAR2(50);
    TYPE my_cursor IS REF CURSOR;
```

```

    the_cursor my_cursor;
BEGIN
    OPEN the_cursor
        FOR 'SELECT title
            FROM books
            WHERE rating = :in_rating'
        USING in_rating;
    DBMS_OUTPUT.PUT_LINE('All books with a rating of ' || in_rating || ':');
    LOOP
        FETCH the_cursor INTO matching_title;
        EXIT WHEN the_cursor%NOTFOUND;
        DBMS_OUTPUT.PUT_LINE(matching_title);
    END LOOP;
    CLOSE the_cursor;
END list_a_rating;

```

This procedure uses a cursor (named `the_cursor`) to return only rows where the book has the specified rating (`in_rating` parameter), and uses a loop to output the title of each book with that rating.

Click the Save icon to save the procedure.

As a usage example, after creating the procedure named `LIST_A_RATING`, if you have inserted data into the `BOOKS` table (for example, using the `INSERT` statements in [Insert Data into the Tables](#) (page 5-7)), you could use the following statement to return all books with a rating of 10:

```
CALL list_a_rating(10);
```

To run this procedure within SQL Developer, right-click `LIST_A_RATING` in the Connections navigator hierarchy display and select **Run**. Under **PL/SQL Block** in the Run PL/SQL dialog box, change `IN_RATING => IN_RATING` to `IN_RATING => 10`, and click **OK**. The Log window display will now include the following output:

```

All books with a rating of 10:
Moby Dick
Software Wizardry

```

5.8 Debug a PL/SQL Procedure

If you want to practice debugging a PL/SQL procedure with SQL Developer, create a procedure that is like the `list_a_rating` procedure that you created in [Create a PL/SQL Procedure](#) (page 5-9), but with a logic error. (The coding is also deliberately inefficient, to allow the display of the rating in a variable.)

Before you can debug the procedure, you must ensure that the user associated with the database connection has the `DEBUG CONNECT SESSION` and `DEBUG ANY PROCEDURE` privileges.

To create this procedure, if you are not already connected, connect to the database as the user for the schema you are using for this tutorial. Right-click the Procedures node in the schema hierarchy on the left side, select **New Procedure**, and enter the following information using the Create PL/SQL Procedure dialog box.

Object Name: `list_a_rating2`

Click **OK**. A source window for the new procedure is opened. Enter (or copy and paste) the following procedure text, replacing any existing text:

```

CREATE OR REPLACE
PROCEDURE list_a_rating2(in_rating IN NUMBER) AS

```

```

matching_title VARCHAR2(50);
matching_rating NUMBER;
TYPE my_cursor IS REF CURSOR;
the_cursor my_cursor;
rating_cursor my_cursor;
BEGIN
  OPEN the_cursor
    FOR 'SELECT title
        FROM books
        WHERE rating <= :in_rating'
    USING in_rating;
  OPEN rating_cursor FOR 'SELECT rating FROM books WHERE
    rating <= :in_rating' USING in_rating;
  DBMS_OUTPUT.PUT_LINE('All books with a rating of ' || in_rating || ':');
  LOOP
    FETCH the_cursor INTO matching_title;
    FETCH rating_cursor INTO matching_rating;
    EXIT WHEN the_cursor%NOTFOUND;
    DBMS_OUTPUT.PUT_LINE(matching_title);
  END LOOP;
  CLOSE the_cursor;
  CLOSE rating_cursor;
END list_a_rating2;

```

This procedure contains a logic error in the definition of the `the_cursor`: it selects titles where the rating is less than or equal to a specified rating, whereas it should select titles only where the rating is equal to the specified rating.

Click the Save icon to save the procedure.

Assume that you wanted to run this procedure and list all books with a rating of 10. Right-click `LIST_A_RATING2` in the Connections navigator hierarchy display and select **Run**. Under **PL/SQL Block** in the Run PL/SQL dialog box, change, change `IN_RATING => IN_RATING` to `IN_RATING => 10`, and click **OK**. In the Log window, however, you see unexpected output: many titles are listed, including some with ratings other than 10. So, you decide to debug the procedure.

To debug the procedure, follow these steps:

1. Click the Compile for Debug icon in the toolbar under the `LIST_A_RATING2` tab.
2. Set two breakpoints by clicking in the left margin (left of the thin vertical line) beside each of these two lines:

```

FETCH the_cursor INTO matching_title;
FETCH rating_cursor INTO matching_rating;

```

Clicking in the left margin toggles the setting and unsetting of breakpoints. Clicking beside these two lines will enable you to see the values of the `matching_title` and `matching_rating` variables as execution proceeds in debug mode.

3. Click the Debug icon, and in the Run PL/SQL dialog box change `IN_RATING => IN_RATING` to `IN_RATING => 10`; then click **OK**
4. Click **View**, then **Debugger**, then **Data** to display the Data pane. (Tip: Expand the Name column width so that you can see `MATCHING_RATING`.)
5. Press the **F9** key (or click **Debug**, then **Resume**) to have execution proceed, stopping at the next breakpoint.

6. Repeatedly press the **F9** key (or click **Debug**, then **Resume**), noticing especially the value of `MATCHING_RATING` as each row is processed. You will notice the first incorrect result when you see that the title *Get Rich Really Fast* is included, even though its rating is only 1 (obviously less than 10).
7. When you have enough information to fix the problem, you can click the Terminate icon in the debugging toolbar.

From this debugging session, you know that to fix the logic error, you should change `rating <= :in_rating` to `rating = :in_rating` in the definition of `the_cursor`.

Related Topics:

[Running and Debugging Functions and Procedures](#) (page 1-36)

[Run/Debug/Profile PL/SQL](#) (page 6-132)

5.9 Use the SQL Worksheet for Queries

You can use the SQL Worksheet to test SQL statements using a database connection. To display the worksheet, from the **Tools** menu, select **SQL Worksheet**. In the Select Connection dialog box, select the database connection that you used to create the `BOOKS`, `PATRONS`, and `TRANSACTIONS` tables for the tutorial.

In the **Enter SQL Statement** box, enter the following statement (the semicolon is optional for the SQL Worksheet):

```
SELECT author_last_name, title FROM books;
```

Notice the automatic highlighting of SQL keywords (`SELECT` and `FROM` in this example).

Click the Execute SQL Statement icon in the SQL Worksheet toolbar. The results of the query are displayed on the **Results** tab under the area in which you entered the SQL statement.

In the **Enter SQL Statement** box, enter (or copy and paste) the following statement, which is the same as the `SELECT` statement in the view you created in Create a View:

```
SELECT p.patron_id,
       p.last_name,
       p.first_name,
       t.transaction_type,
       t.transaction_date
FROM patrons p, transactions t
WHERE p.patron_id = t.patron_id
ORDER BY p.patron_id, t.transaction_type;
```

Click the Execute SQL Statement icon in the SQL Worksheet toolbar, and view the results of the query.

Click the Execute Explain Plan icon in the SQL Worksheet toolbar to see the execution plan (displayed on the Explain tab) that Oracle Database follows to execute the SQL statement. The information includes the optimizer strategy and the cost of executing the statement.

See Also:

For information about the SQL Worksheet user interface, see [Using the SQL Worksheet](#) (page 1-41).

For information about how to generate and interpret execution plans, see *Oracle Database SQL Tuning Guide*

5.10 Script for Creating and Using the Library Tutorial Objects

The following statements create and use the database objects that you have created (or will create) for the tutorial in [SQL Developer Tutorial: Creating Objects for a Small Database](#) (page 5-1). You can view these commands to help you understand the library database objects that are covered in the tutorial.

```
-- Clean up from any previous tutorial actions.
DROP TABLE transactions;
DROP TABLE books;
DROP TABLE patrons;
DROP SEQUENCE patron_id_seq;
DROP SEQUENCE transactions_seq;
DROP TRIGGER transactions_trg;
DROP VIEW patrons_trans_view;
DROP PROCEDURE list_a_rating;
DROP PROCEDURE list_a_rating2;

set serveroutput on

-- Create objects.

CREATE TABLE books (
  book_id VARCHAR2(20),
  title VARCHAR2(50)
    CONSTRAINT title_not_null NOT NULL,
  author_last_name VARCHAR2(30)
    CONSTRAINT last_name_not_null NOT NULL,
  author_first_name VARCHAR2(30),
  rating NUMBER,
  CONSTRAINT books_pk PRIMARY KEY (book_id),
  CONSTRAINT rating_1_to_10 CHECK (rating IS NULL OR
    (rating >= 1 and rating <= 10)),
  CONSTRAINT author_title_unique UNIQUE (author_last_name, title));

CREATE TABLE patrons (
  patron_id NUMBER,
  last_name VARCHAR2(30)
    CONSTRAINT patron_last_not_null NOT NULL,
  first_name VARCHAR2(30),
  street_address VARCHAR2(50),
  city_state_zip VARCHAR2(50),
  location MDSYS.SDO_GEOMETRY,
  CONSTRAINT patrons_pk PRIMARY KEY (patron_id));

CREATE TABLE transactions (
  transaction_id NUMBER,
  patron_id CONSTRAINT for_key_patron_id
    REFERENCES patrons(patron_id),
  book_id CONSTRAINT for_key_book_id
    REFERENCES books(book_id),
```

```
transaction_date DATE
    CONSTRAINT tran_date_not_null NOT NULL,
transaction_type NUMBER
    CONSTRAINT tran_type_not_null NOT NULL,
CONSTRAINT transactions_pk PRIMARY KEY (transaction_id));

CREATE SEQUENCE patron_id_seq
    START WITH 100
    INCREMENT BY 1;

-- The sequence for the transaction_id
-- in the tutorial is created automatically,
-- and may have the name TRANSACTIONS_SEQ.
CREATE SEQUENCE transactions_seq
    START WITH 1
    INCREMENT BY 1;

-- The before-insert trigger for transaction ID values
-- in the tutorial is created automatically,
-- and may have the name TRANSACTIONS_TRG.
CREATE OR REPLACE TRIGGER transactions_trg
    BEFORE INSERT ON TRANSACTIONS
    FOR EACH ROW
    BEGIN
        SELECT TRANSACTIONS_SEQ.NEXTVAL INTO :NEW.TRANSACTION_ID FROM DUAL;
    END;
/

CREATE VIEW patrons_trans_view AS
    SELECT p.patron_id,
           p.last_name,
           p.first_name,
           t.transaction_type,
           t.transaction_date
    FROM patrons p, transactions t
    WHERE p.patron_id = t.patron_id
    ORDER BY p.patron_id, t.transaction_type;

-- Procedure: List all books that have a specified rating.
CREATE OR REPLACE PROCEDURE list_a_rating(in_rating IN NUMBER) AS
    matching_title VARCHAR2(50);
    TYPE my_cursor IS REF CURSOR;
    the_cursor my_cursor;
BEGIN
    OPEN the_cursor
        FOR 'SELECT title
            FROM books
            WHERE rating = :in_rating'
        USING in_rating;
    DBMS_OUTPUT.PUT_LINE('All books with a rating of ' || in_rating || ':');
    LOOP
        FETCH the_cursor INTO matching_title;
        EXIT WHEN the_cursor%NOTFOUND;
        DBMS_OUTPUT.PUT_LINE(matching_title);
    END LOOP;
    CLOSE the_cursor;
END;
/
show errors;

-- Insert and query data.
```



```
INSERT INTO books VALUES ('A1111', 'Moby Dick', 'Melville', 'Herman', 10);
INSERT INTO books VALUES ('A2222', 'Get Rich Really Fast', 'Scammer', 'Ima', 1);
INSERT INTO books VALUES ('A3333', 'Finding Inner Peace', 'Blissford', 'Serenity',
null);
INSERT INTO books VALUES ('A4444', 'Great Mystery Stories', 'Whodunit', 'Rodney', 5);
INSERT INTO books VALUES ('A5555', 'Software Wizardry', 'Abugov', 'D.', 10);

INSERT INTO patrons VALUES (patron_id_seq.nextval,
'Smith', 'Jane', '123 Main Street', 'Mytown, MA 01234', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
'Chen', 'William', '16 S. Maple Road', 'Mytown, MA 01234', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
'Fernandez', 'Maria', '502 Harrison Blvd.', 'Sometown, NH 03078', null);
INSERT INTO patrons VALUES (patron_id_seq.nextval,
'Murphy', 'Sam', '57 Main Street', 'Mytown, MA 01234', null);

INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (100, 'A1111', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (100, 'A2222', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (101, 'A3333', SYSDATE, 3);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (101, 'A2222', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (102, 'A3333', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (103, 'A4444', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (100, 'A4444', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (102, 'A2222', SYSDATE, 2);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (102, 'A5555', SYSDATE, 1);
INSERT INTO transactions (patron_id, book_id,
transaction_date, transaction_type)
VALUES (101, 'A2222', SYSDATE, 1);

-- Test the view and the procedure.
SELECT * FROM patrons_trans_view;
CALL list_a_rating(10);
```

SQL Developer Dialog Boxes and Wizards

SQL Developer uses dialog boxes for creating and editing database connections and objects in the database (tables, views, procedures, and so on). The dialog boxes sometimes have multiple tabs, each reflecting a logical grouping of properties for that type of object.

For an explanation of any dialog box or tab, click the **Help** button or press the **F1** key.

The dialog boxes and wizards are not presented here in any rigorous order, because the help for each is an independent piece of information and is normally seen when you click Help or press F1 in that context.

Note:

For all Name fields, any name that you type is automatically converted to and stored in the database metadata in uppercase, unless you enclose the name in quotation marks (" "). (Names of database objects in SQL and PL/SQL statements are not case-sensitive.)

To include lowercase characters, special characters, or spaces in object names, enclose the name in quotation marks (" ") when you type it. Example: "My table"

Related Topics:

[SQL Developer Concepts and Usage](#) (page 1-1)

[Database Objects](#) (page 1-13)

6.1 Add Extension

This dialog box is displayed when you click Add in the File Types pane of SQL Developer preferences.

Extension: Specify the file extension, including the period (for example, .xyz).

After you click OK, you can select that extension and modify its details, including the file type, content type, and whether to have files with the extension automatically opened by SQL Developer.

Related Topics:

[File Types](#) (page 1-111)

[SQL Developer Preferences](#) (page 1-92)

6.2 Add Schema Error

This dialog box is displayed when you specify an invalid file after clicking Add in the XML Schemas pane of SQL Developer preferences. A list of the errors is displayed.

Ignore errors and continue registration: Ignores the errors and allows the process of registering the specified XML schema or schemas after you click **OK**. Choose this option only if you are sure that the apparent errors can be ignored.

Skip registration and open using SQL Developer: Opens an editing window in which you can correct any errors.

Related Topics:

[XML Schemas](#) (page 1-122)

[SQL Developer Preferences](#) (page 1-92)

6.3 Advanced Properties (Connections)

This dialog box is displayed if you click Advanced in the [Create/Edit/Select Database Connection](#) (page 6-11) dialog box.

6.3.1 Proxy tab

This information applies to proxy connections. For an explanation of proxy authentication, see [Connections with Proxy Authentication](#) (page 1-32).

Proxy Type: *User Name* for authentication by proxy user name and password, or *Distinguished Name* for authentication by proxy user name and distinguished name.

Proxy User: Name of the user to be used for authentication for this connection.

Proxy Password (if Proxy Type is User Name): Password for the specified proxy user.

Distinguished Name (if Proxy Type is Distinguished Name): Distinguished name for the specified proxy user.

6.4 Application Migration

The Application Wizard enables you to migrate a DBLib or CTLib Sybase application to Oracle by specifying a directory with source code from the DBLib or CTLib Sybase application. You will be able to see what changes must be made to migrate the application to Oracle, and to perform the migration.

The Welcome page includes a description of the wizard's operation and options. If you want to continue to see this Welcome page in the future, do *not* enable (do not check) the **Do not show this page again** option.

Follow the instructions on each page of the wizard.

On the Overview page, specify the **Application Name** and optionally a **Description** of the application.

On the File Selection page, specify the application root directory and other information to determine the files to be added to the project:

Application Directory: Root directory for the application. All files and folders under this directory will be scanned and added to the project.

Database Type: Type of third-party database for the migration (for example, *sybase*).

Output Actions: What the wizard should do with the output (for example, *Output to directory*, and then specify the **Output Directory**).

Use Custom Rules: Whether to use custom rules for the output; and if so, the **Rules Directory** for the custom rules.

Related Topics

- [SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

6.5 Associate Repository

This dialog box is displayed if you click **Tools**, then **Migration**, then **Repository Management**, then **Associate Repository**.

Associate Repository: Name of the database connection to use to create a migration repository. The objects associated with the migration repository are created in the schema of the user associated with the selected connection.

6.6 Cart Error (Objects Not Available)

This box, with the text *Some objects included in the cart are not available. Review cart contents.*, is displayed when you attempt to perform a **Cart** operation when one or more of the database objects in the selected cart are no longer available. For example, if you attempt to export some tables that include *TABLE_1*, but *TABLE_1* was dropped or the database connection was deleted after that table was added to the cart, then this error box is displayed.

To fix the problem, ensure that you are attempting to perform the **Cart** operation only on database objects that exist and are available.

Related Topics

- [Using the Cart](#) (page 1-68)

6.7 Change Type

Use this dialog box to change the data type of a column in a captured model before you perform the migration.

Source Data Type: Specify the new data type for the column.

Any remaining fields in the dialog box depend on the **Source Data Type** that is selected.

6.8 Check for Updates

When you click **Help** and then **Check for Updates**, you can check for and download available SQL Developer updates. The following pages may be displayed. (If you have enabled the SQL Developer preference to check for updates automatically at startup, and if you click to see available updates at startup, the **Updates** page is displayed.)

If you are unable to check for updates because your system is behind a firewall, you may need to set the SQL Developer user preferences for [Web Browser and Proxy](#) (page 1-121).

1. **Source:** Select the source or sources to be checked for available updates: any or all of some specified online update centers, or a local ZIP file containing an update bundle. You can also click **Add** to add a user-defined update center.
2. **Updates:** If any updates are available from the selected source or sources, select those that you want to download. The available updates include certain third-party JDBC drivers, which require that you agree to the terms of their licenses.

The **Show Upgrades Only** option restricts the display to upgrades of currently installed SQL Developer components. To enable the display of all new and updated components, whether currently installed or not, uncheck this option.

After you click Next, you may be prompted to enter your Oracle Web Account user name and password. If you do not have an account, you can click the Sign Up link.
3. **License Agreements** (displayed only if you selected any updates that require a license agreement): For each update that requires you to agree to the terms of a license, review the license text and click **I Agree**. You must do this for each applicable license.
4. **Download:** If you selected any updates to download, this page displays the progress of the download operation.
5. **Summary:** Displays information about the updates that were downloaded. After you click Finish, you will be asked if you want to install the updates now and restart SQL Developer.

6.9 Choose Directory

This is a standard box for choosing a directory in which to place files: use **Location** to navigate to (double-clicking) the folder in which to save the files, or enter a directory name. If the directory does not already exist, it is created.

6.10 Clone PDB to Oracle Cloud

Use this dialog box to clone a PDB to the Oracle Cloud. It unplugs the source PDB, plugs it into the Cloud CDB destination, then either plugs the source PDB back into its CDB or deletes it from its CDB.

Source PDB: Name of the PDB to be cloned to the Oracle Cloud.

Destination Connection: Name of the multitenant container database (CDB) connection in which to plug in the cloned PDB.

Action after clone: Determines what happens to the original (source) PDB after its clone is unplugged from its current CDB and plugged into the CDB at the destination connection.

- **RePlug:** The source PDB is plugged back into its CDB. (Note that after the source PDB is replugged, it and the clone made at the destination connection are separate and independent; you can modify metadata and data in one without affecting the other, if you wish.)
- **Delete:** The source PDB is deleted from its CDB.

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.11 Clone Pluggable Database

Use this dialog box to clone a PDB. It creates a new PDB from (using basic specifications of) a specified PDB.

Database Name: Name of the new PDB.

Source PDB: Source PDB for the cloning operation.

Database Link: If the source PDB is in a remote CDB, specify the name of the database link to use for connecting to the remote CDB.

Storage: You can specify storage limits for the PDB total size or temporary tablespace usage, or both; or unlimited storage for either or both.

File Name Conversions: Determines how the database generates the names of files (such as data files and wallet files) for the PDB.

- **None:** The database first attempts to use Oracle Managed Files to generate file names. If you are not using Oracle Managed Files, then the database uses the `PDB_FILE_NAME_CONVERT` initialization parameter to generate file names.
- **Custom Names:** Select a Source Files/Target Files pair.
- **Custom Expressions:** Specify one or more Source File Expression/Target File Expression pairs. Each pair item is a string found in names of files associated with the seed (when creating a PDB by using the seed), associated with the source PDB (when cloning a PDB), or listed in the XML file (when plugging a PDB into a CDB).

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.12 Color Palette and Custom Colors

You can use the color palette editor to select a color from the supplied Available Colors or saved Custom Colors.

You can also create a color by using a gradient box or by specifying the RGB (Red, Green, Blue) values for the color.

6.13 Configure Extension

This dialog box, which is displayed if you click Configure for Versioning Support in the Extensions preferences pane, enables you to select from among available versioning support extensions for SQL Developer. For information about using versioning with SQL Developer, see [Using Versioning](#) (page 1-53).

If you change any existing settings, you will need to restart SQL Developer.

6.14 Configure File Type Associations

This dialog box, which is displayed the first time you start SQL Developer, enables you to associate certain file types with SQL Developer. If a file type is associated with SQL Developer, files with that type's extension will automatically be opened by SQL Developer when you double-click the file name. Any previous association for that file type is replaced.

If you do not associate a file type with SQL Developer, any existing association for that file is unchanged.

After you close this box, you can change the associations for these file types and many others by clicking **Tools** and then **Preferences**, and selecting **File Types** (see [File Types](#) (page 1-111)).

6.15 Copy Columns

This dialog box is displayed if you click the Copy Columns icon when specifying column definitions for a table.

Schema: The schema that owns the table from which to copy column definitions.

Table: The table within the selected schema.

Columns: A list of the columns in the table. Select one or more columns to be copied into the table that you are creating or editing, and click **OK**.

6.16 Copy Objects

This dialog box is displayed if you click the Copy Objects icon in the Cart window.

Open Configuration icon: Opens a previously saved XML configuration file, to use its settings as defaults for this use of the dialog box.

Save Configuration icon: Saves the current settings in the dialog box to an XML file, which you can later open to use for Cart operations of this type.

Destination Connection: Database connection into which to copy the objects.

Copy DDL: Copies the object definitions. For destination objects with the same names as source objects of the same type (for example, if both contain a table named EMPLOYEES), specify whether to not perform the copy (that is, do *not* replace the destination objects) or to perform the copy (that is, replace the existing destination objects with the source objects).

Copy Data: Copies the data for any tables and views that are copied. If you do not select this option, any copied tables or views are empty in the destination connection.

Truncate Destination Data Before Copying: If a table or view of the same name already exists in the destination connection (for example, if an EMPLOYEES table exists in the source and destination), this option deletes any existing data in the destination object before copying the data from the source connection. If this option is not selected, the copied data is appended to the existing data in the destination object.

Related Topics:

[Using the Cart](#) (page 1-68)

[Save or Open Cart Tool Configuration](#) (page 6-134)

[Database Connections](#) (page 1-28)

6.17 Copy to Hadoop / Append to Hadoop Table

Use this wizard to copy a table from Oracle Database to HDFS and create a Hive table, or to append data to an existing Hive table that was created using Copy to Hadoop. To use the wizard, you must understand the information in [Apache Hadoop Connectors Support in SQL Developer](#) (page 1-79).

The copy operation creates Oracle Data Pump files from an Oracle Database table, and copies them to HDFS. It also creates a Hive external table over the Data Pump files. It can optionally convert the Oracle Data Pump files to Parquet or ORC format.

You can initiate Copy by right-clicking the Tables icon in a Hive database schema.

You can initiate Append to an Existing Hive External Table by right-clicking the icon for that Hive table.

The wizard steps specify details for the source and destination objects, and selection criteria for the source data content to be copied.

Copy Method

Select Direct Copy to Staging External Table method.

The Direct Copy method copies data from an Oracle Database table directly to HDFS. It copies data from one Oracle table, rows can be selected by a WHERE clause filter. Direct Copy requires an SSH connection to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. The Direct Copy method is recommended.

The Staging External Table method exports data to a file system location on the database system, and then copies them to HDFS. The Staging External Table method requires an SSH connection to the database system, connecting as the user running the database process (typically “oracle”).

SSH Connection: You must select an SSH connection.

For the Direct Copy method, this is an SSH connection to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. As discussed in the [Apache Hadoop Connectors Support in SQL Developer](#) (page 1-79) section, the `.sqldev_cp2hadoop_env` file must have been created in the home directory of this SSH connection.

For the Staging External Table method, this is an SSH connection to the database system.

The Copy to Hadoop job will run from the node of the SSH connection.

Source and Destination

Oracle Database Source: Specify the connection, database object, Oracle directory, Hadoop cluster (if Oracle Big Data SQL is configured to access multiple clusters), and degree of parallelism.

Hive Destination: Specify the Hive external table name and HDFS root directory for the data files.

Source Selection Criteria

Enter the selection criteria for the source table. The selection criteria is used to specify table rows to copy in the copy operation.

Optionally, select the columns to include in the copy operation.

Optionally, select parameters and specify a WHERE clause.

Refresh: Click to view sample results.

Source Connection

Specify Wallet credentials if using Oracle Wallet. These are credentials to access the table you are copying from.

Select the **Delete Deployment Files** checkbox on the SSH host to clean up the files for this job. For debugging, deselect the checkbox.

Summary

You can review the options that you specified before clicking **Finish** to submit the job. Click the Process tab to view the steps that will be executed as a part of the job.

6.18 Copy to Oracle

This dialog box is displayed if you try to drag and drop an object from one connection to another connection. The options shown depend on the type of object. For copying a table, the options include the following.

Copy DDL: Copies the object definitions. For existing destination objects with the same names as source objects of the same type (for example, if both contain a table named EMPLOYEES), specify either **Do Not Replace** (that is, do *not* perform the copy and replace the destination objects) or **Replace** (that is, replace the existing destination objects with the source objects).

Copy Data: Copies the data for any tables that are copied. If you do not select this option, any copied tables are empty in the destination connection. For existing destination objects with the same names as source objects of the same type (for example, if both contain a table named EMPLOYEES), specify one of the following:

- **Do Not Copy Data:** Do not copy data from the source object.
- **Append:** Append the copied data in the destination object.
- **Truncate Destination Data Before Copying:** Delete any existing data in the destination object before copying the data from the source connection.

Related Topics

- [Database Connections](#) (page 1-28)
- [Database Objects](#) (page 1-13)

6.19 Component Palette

The Component Palette displays the elements that you can drag and drop. The components available for selection in the palette vary depending on the content of the active editor window.

To insert a component into a file open in the active editor, drag the component from the palette to an insertion point in the editor. In some file types you can click a component in the palette and then click in the editor to insert the component.

Enter the appropriate information in the dialog box that is displayed (for example, Insert CDATA or Insert Processing Instruction in an XML file).

6.20 Component Palette: Configure Component Palette

Lets you configure the component palette. Note that some page types cannot be edited or removed, and most existing component types cannot be added to, edited, or removed.

Add: Displays the [Component Palette: Create Palette Page](#) (page 6-9) dialog box.

Remove: Deletes the selected page from the palette.

Rename: Renames a specified page.

6.21 Component Palette: Create Palette Page

Lets you create a new page for the component palette. Specify a name of the page, and select the type of page from a list.

Page Name: Name of the page. Suggestion: Include the type of page in the name, perhaps naming pages in the form *name_type_page*.

Page Type: Page type, selected from the list.

6.22 Component Palette: New/Edit Code Snippet

Lets you create or edit a code snippet, which you will be able to drag into files that you edit.

Name: Name for the code snippet.

Image: Icon image to be associated with the snippet.

Code: Code for the snippet.

6.23 Component Palette: New Section/Rename Section

Lets you create or rename a section in the Code Snippets panel of the Component Palette.

You can create sections in the Code Snippets panel to organize your snippets better. For example, if you have a group of code snippets that pertain to mathematical functions, you can create a new section called Math and group the related snippets under it.

Name: Name of the section.

6.24 Confirm Drop Application

This dialog box is displayed when you right-click an Application Express application and select Drop. To drop the application, click **Yes**; to keep (not drop) the application, click **No**.

If the application contains an uninstall script, that script is run before the application is dropped.

Related Topics

- [Applications \(Application Express 3.0.1 and Later\)](#) (page 1-14)

6.25 Confirm Running SQL

This dialog box is displayed in certain situations when SQL Developer needs to run a setup script on the server. The script is displayed in a text box, where you can view or edit the contents. To allow the script to run, click **Yes**; to prevent the script from running, click **No**.

6.26 Connection Has Uncommitted Changes

This dialog box is displayed if you try to end the active database session while there are transactions to be committed. Select the appropriate option and click **OK**.

To commit the changes and end the session, select **Commit Changes**. To roll back the changes and end the session, select **Rollback Changes**. To cancel the attempt to end the session, select **Abort Connection Disconnect**. (Selecting Abort Connection Disconnect and clicking OK has the same effect as clicking Cancel.)

6.27 Create New Object

This dialog box is displayed if you click File, then New. Specify the type of object to create. After you click OK, the dialog box for creating that type of object is displayed.

Related Topics

- [Database Objects](#) (page 1-13)

6.28 Create/Edit Chain

This dialog box or pane is used for creating or editing an Oracle Scheduler chain.

(To create a chain, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_CHAIN procedure.)

Name: Name of the chain.

Enabled: Enables the chain. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the chain is created.)

Description: Optional text string that can be used to describe the chain.

Related Topics

- [Chains](#) (page 1-14)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.29 Create/Edit Credential

This dialog box or pane is used for creating a new Oracle Scheduler credential or editing an existing credential.

(To create a credential, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_CREDENTIAL procedure.)

Name: Name of the credential. Cannot be set to NULL. It is converted to upper case unless enclosed in double-quotes. For an existing credential, this field is read-only; to change the name, you must drop the credential and create a new credential with the desired name.

Enabled: If this option is selected, the credential is enabled; if this option is not enabled, the credential is not enabled.

Description: Optional text string that can be used to describe the credential.

User Name: User name to use for login to the host operating system or remote Oracle database. This cannot be set to NULL and is case sensitive. It cannot contain double quotes or spaces.

Password: Password for the user name. This cannot be set to NULL and is case sensitive. The password is stored obfuscated and is not displayed in the Scheduler dictionary views.

Database Role: (Reserved for future use.)

Windows Domain: For a Windows remote executable target, this is the domain that the specified user belongs to. The domain is converted to uppercase automatically.

Related Topics

- [Credentials](#) (page 1-15)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.30 Create/Edit/Select Database Connection

The database connection dialog box displays any existing connections. Depending on the context, you can select a connection to connect to the database, edit the information about existing connections, or specify information while creating a new connection. (See [Creating and Editing Connections](#) (page 6-17).)

Connection Name: An alias for a connection to the database using the information that you enter. (The connection name is not stored in the database, and the connection is not a database object.) Suggestion: Include the database name (SID) and user name in the connection name. Example: personnel_herman for connecting to the personnel database as user Herman.

Username: Name of the database user for the connection. This user must have sufficient privileges to perform the tasks that you want to perform while connected to the database, such as creating, editing, and deleting tables, views, and other objects.

Password: Password associated with the specified database user.

Save Password: If this option is checked, the password is saved with the connection information, and you will not be prompted for the password on subsequent attempts to connect using this connection.

Connection Color: Lets you specify a standard or custom color for the border of any SQL Worksheet and other windows associated with the connection. The border can be a helpful visual indicator for preventing confusion; for example, if you use red borders for connections to the production database, it might prevent you from mistakenly dropping a table in the production environment when you intended to drop it in the development environment.

Connection Color does not affect the display of the connection name in the Connections navigator.

Information for Database-Specific Tabs:

- [Oracle tab](#) (page 6-12)
- [TimesTen tab](#) (page 6-14)
- [Amazon Redshift tab](#) (page 6-14)
- [DB2 tab](#) (page 6-15)
- [Hive tab](#) (page 6-15)
- [JDBC tab](#) (page 6-16)
- [MySQL tab](#) (page 6-16)
- [PostgreSQL tab](#) (page 6-16)
- [SQL Server and Sybase tab](#) (page 6-17)

- [Teradata tab](#) (page 6-17)

Related Topics

- [Database Connections](#) (page 1-28)

6.30.1 Oracle tab

The following information applies to a connection to an Oracle Database.

Connection Type: Select Basic, TNS, LDAP (Lightweight Directory Access Protocol), Advanced, Local/Bequeath (using *bequeath* protocol: if the client and database exist on the same computer, then a client connection can be passed directly to a dedicated server process without going through the listener), SSH, or Cloud PDB. (The display of fields changes to reflect any change in connection type.)

Role: The set of privileges to be associated with the connection. For a user that has been granted the SYSDBA system privilege, you can specify a connection that includes the privilege.

OS Authentication: If this option is checked, control of user authentication is passed to the operating system (OS). This allows the specified user to connect to the database by authenticating that user's OS username in the database. No password is associated with the connection since it is assumed that OS authentication is sufficient.

Kerberos Authentication: If this option is checked, credentials can be shared across many Kerberos-enabled applications (for example, to have the same username and password for both the operating system and Oracle Database). Thick driver configuration is done through `sqlnet.ora` (`sqlnet.authentication_services=(KERBEROS)` and related parameters), so no username and password are needed. Thin driver configuration uses the configuration (`.conf`) file and the credentials cache, and uses a service principal and password.

Proxy Connection: If this option is checked, proxy authentication will be used. Displays the [Advanced Properties \(Connections\)](#) (page 6-2) dialog box.

Basic Connection Type

Hostname: Host system for the Oracle database.

Port: Listener port.

SID: Database name.

Service Name: Network service name of the database (for a remote database connection over a secure connection).

TNS Connection Type

Network Alias: Oracle Net alias for the database. (The list for selecting a network alias is initially filled from all `tnsnames.*` files on your system, unless you have set the Database: Advanced preference **Tnsnames Directory** to identify the location of the `tnsnames.ora` file to be used.)

Connect Identifier: Oracle Net connect identifier.

LDAP Connection Type

Enterprise users are authenticated with the Lightweight Directory Access Protocol (LDAP) server. The user login information must be configured in the LDAP server and mapped to a schema in the database. Support for LDAP-compliant directory servers provides a centralized vehicle for managing and configuring a distributed Oracle

network. The directory server can replace client-side and server-side localized tnsnames.ora files.

LDAP Server: Select from the list (from <DIRECTORY_SERVER> entries in the ldap.ora file); or enter the directory server location and port (either SSL or non-SSL), for example: `system123.example.com:389:636` (ldap-system:nonssl-port:ssl-port)

Context: LDAP administrative context. The contexts available in the selected server are listed.

DB Service: Database connection information: click **Load** to display a list of database services associated with the selected context. (If an error is displayed, no database services are associated with this context.) If a connection uses the OCI/Thick driver (see the **Use OCI/Thick** preference under Database: Advanced), the system on which SQL Developer is running must have an Oracle Client installation that contains the JDBC and orai18n libraries, these libraries must be present on the path, and the Oracle Client installation must be version 10.2 or later.

Advanced Connection Type

Custom JDBC URL: URL for connecting directly from Java to the database; overrides any other connection type specification. If you are using TNS or a naming service with the OCI driver, you must specify this information: Example:

```
jdbc:oracle:thin:scott/@localhost:1521:orcl
```

Note that in this example, the "/" is required, and the user will be prompted to enter the password.

To use a custom JDBC URL, the system on which SQL Developer is running must have an Oracle Client installation that contains the JDBC and orai18n libraries, is present on the path, and is version 10.2 or later.

SSH Connection Type

Port Forward: Name of the port forward for the SSH session.

SID: Database name.

Service Name: Network service name of the database.

Cloud PDB Connection Type

For Oracle Database Exadata Express Cloud Service users. See the "Connecting SQL Developer" section in *Oracle Cloud Using Oracle Database Exadata Express Cloud Service* for details.

See Also:

For Kerberos authentication options, see the [Database](#) (page 1-99): Advanced preferences.

For configuring Kerberos authentication, see *Oracle Database Advanced Security Guide*

[Connections with Proxy Authentication](#) (page 1-32)

6.30.2 TimesTen tab

The following information applies to a connection to an Oracle TimesTen In-Memory Database.

For **Username** and **Password**, specify the user name and password of the user account in the TimesTen database.

DSN: Data source name. Select an existing DSN (if any are displayed), or *User-specified* to create a new DSN. A DSN is a character string that identifies a TimesTen database and includes connection attributes to be used when connecting to the database. A DSN has the following characteristics: its maximum length is 32 characters; it cannot contain spaces; and it consists of ASCII characters except for the following: [] { } , ; ? * = ! @ \

Connection Type (if DNS is user-specified): *C/S* for client/server mode or *Direct* for direct mode

Connection String: Connection attributes including database attributes, first connection attributes, general connection attributes, NLS attributes, and Cache Connect attributes.

Oracle Password (for Cache): The password for the TimesTen user account on the Oracle Database. (See the TimesTen documentation for more information.)

See Also:

For information about connection attributes, see *Oracle TimesTen Application-Tier Database Cache User's Guide*.

For more information about SQL Developer support for TimesTen, see [Oracle TimesTen In-Memory Database Support](#) (page 1-125).

For detailed usage and reference information about Oracle TimesTen, see the online documentation that is included in the TimesTen installation. For additional information, go to: <http://www.oracle.com/technetwork/database/timesten/>

6.30.3 Amazon Redshift tab

The following information applies to a connection to an Amazon Redshift database.

Note that the Amazon Redshift tab is not displayed unless you first download the appropriate connection driver, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for Database: Third Party JDBC Drivers to add the driver.

Amazon Redshift JDBC URL: The JDBC URL of the cluster has the following format: `jdbc:redshift://endpoint:port/database`, where *endpoint* is the endpoint of the Amazon Redshift cluster, *port* is the port number that was specified when the cluster was launched, and *database* is the database created for the cluster.

You can obtain the URL for the JDBC connection to the database from the cluster configuration in the Amazon Web Services console.

See Also:

For more details about obtaining the JDBC URL, see <http://docs.aws.amazon.com/redshift/latest/mgmt/configure-jdbc-connection.html#obtain-jdbc-url>

6.30.4 DB2 tab

The following information applies to a connection to an IBM DB2 database.

Note that to connect to an IBM DB2 database, you must first download the db2jcc.jar and db2jcc_license_cu.jar files, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for Database: Third Party JDBC Drivers to add these files.

Platform: UDB.

Host Name: Host system for the IBM DB2 database.

Port: TCP/IP Port on which the IBM DB2 server will listen.

Enter Database: Name of the IBM DB2 database.

6.30.5 Hive tab

Click the Hive tab (next to the Oracle tab) to enter Hive connection details.

The following information applies to a connection to a Hive database.

Note that to connect to a Hive database, you must first download the Cloudera Hive JDBC driver, which contains the following files: hive-metastore.jar, hive-service.jar, HiveJDBC4.jar, libfb303-0.9.0.jar, libthrift-0.9.0.jar, log4j-1.2.14.jar, ql.jar, slf4j-api-1.5.8.jar, slf4j-log4j12-1.5.8.jar, TCLIServiceClient.jar. (For files whose names indicate a version number, a more recent version might be available.). Note that in SQL Developer 4.2, you need to use the files in Cloudera_HiveJDBC4_*.zip (not Cloudera_HiveJDBC41_*.zip or Cloudera_HiveJDBC3_*.zip). Extract the files from the zip file in a temporary directory. Then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for Database: Third Party JDBC Drivers to add these files.

Username: User name to connect to Hive Server 2. For example, the user name is "oracle" on Oracle Big Data Appliance without Kerberos. **Password** might not be required, depending on how authentication was set up for the database.

Host Name: Host system running Hive Server 2.

Port: TCP/IP Port on which the Hive server will listen (default is 10000).

Database: Name of the Hive database. (default is the name of the base database on Hive. You can create other databases using the CREATE DATABASE command, described at <https://cwiki.apache.org/confluence/display/Hive/LanguageManual+DDL#LanguageManualDDL-Create/Drop/AlterDatabase>.)

To add new parameters for the connection, click **Add**. For example, when you connect to a Kerberos-enabled cluster, add the following parameters:

- **KrbServiceName:** Hive Server 2 service name.
- **AuthMech:** 1 indicates a Kerberos connection.
- **KrbHostFQDN:** Hive Server 2 host name.

- **KrbRealm:** Kerberos realm.

See Also:

For more information about installing Hive JDBC drivers, see “Using Oracle SQL Developer to Connect to Hive” in *Oracle Big Data SQL User’s Guide*.

For information about the user preference pane for third-party JDBC drivers, see [Database](#) (page 1-99).

6.30.6 JDBC tab

The following information applies to a JDBC connection.

JDBC-ODBC Bridge or Other Third Party Driver: Indicates a JDBC to ODBC bridge driver or another third-party driver.

Data Source (JDBC-ODBC Bridge): Name of an existing ODBC data source.

Extra Parameters (JDBC-ODBC Bridge): Additional parameters for the connection.

JDBC URL (Other Third Party Driver): URL for connecting directly from Java to the database; overrides any other connection type specification.

Driver Class (Other Third Party Driver): The name of the driver class that will be used for the connection (for example, `com.microsoft.jdbc.sqlserver.SQLServerDriver`). This name can be found in the JDBC driver specification (usually shipped with the driver).

6.30.7 MySQL tab

The following information applies to a connection to a MySQL database.

Note that to connect to a MySQL database, you must first download the appropriate MySQL connection driver, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for [Database](#) (page 1-99): Third Party JDBC Drivers to add the driver.

Host Name: Host system for the MySQL database.

Port: TCP/IP Port on which the MySQL server will listen.

Choose Database: Name of the MySQL database.

Zero Date Handling: Because the MySQL JDBC driver cannot handle the default 0000-00-00 date, specify one of the following options for handling this date: **Set to NULL** to set it to a null value, or **Round to 0001-01-01** to set it to 0001-01-01.

6.30.8 PostgreSQL tab

The following information applies to a connection to a PostgreSQL database.

For PostgreSQL migrations, only online capture and online data move are supported. (DDL can be generated, but the data must be moved using the online method.)

Note that to connect to a PostgreSQL database, you must first download the appropriate connection driver, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for [Database](#) (page 1-99): Third Party JDBC Drivers to add the driver.

Host Name: Host system for the PostgreSQL database.

Port: TCP/IP Port on which PostgreSQL will listen.

Choose Database: Name of the PostgreSQL database.

6.30.9 SQL Server and Sybase tab

The following information applies to a connection to a Microsoft SQL Server or Sybase Adaptive Server database.

Note that to connect to a Microsoft SQL Server or Sybase Adaptive Server database, you must first download the appropriate connection driver, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for [Database](#) (page 1-99): Third Party JDBC Drivers to add the driver.

Use Default Password: If this option is checked, the SQL Server or Sybase default password for new users and for resetting passwords is used.

Use Windows Authentication: If this option is checked, control of user authentication is passed to the Microsoft Windows operating system. This allows the specified user to connect to the database by authenticating that user's Windows username in the database. No password is associated with the connection since it is assumed that Windows authentication is sufficient.

Host Name: Host system for the Microsoft SQL Server or Sybase Adaptive Server database.

Port: TCP/IP Port on which Microsoft SQL Server or Sybase Adaptive Server will listen.

Retrieve Database: Name of the Microsoft SQL Server or Sybase Adaptive Server database.

6.30.10 Teradata tab

The following information applies to a connection to a Teradata database.

Note that to connect to a Teradata database, you must first download the `tdgssconfig.jar` and a `terajdbc4.jar` files, and then click **Tools**, then **Preferences**, and use the SQL Developer user preference pane for [Database](#) (page 1-99): Third Party JDBC Drivers to add these files. (See also the `readme.txt` file that is included with the `tdgssconfig.jar` and a `terajdbc4.jar` files.)

Host Name: Host system for the Teradata database.

DBS Port: TCP/IP Port on which the Teradata server will listen.

Charset: Character set for the data.

TMODE: Transaction mode: ANSI, TERA (Teradata), or DEFAULT.

To add a connection parameter to the list in the box, click **Add**; to delete a connection parameter from the list, click **Delete**.

6.30.11 Creating and Editing Connections

*To create a new connection when no connections exist, enter the connection information and click **Connect**. To test the connection before you create it, click **Test**.*

*To create a new connection when one or more connections already exist, click to select an existing connection, change the Connection Name to the desired name, edit other connection information as needed, and click **Save** or **Connect** to create the new connection. To test the connection before you create it, click **Test**.*

To edit an existing connection, click in its entry in the Connection Name column, change any connection information except the connection name, and click **Save** or **Connect**. To test the connection before you save changes to it, click **Test**.

6.31 New/Edit Cloud Connection

Use this dialog box to create or edit a connection to an Oracle Cloud Database Schema Service instance, which does not offer Oracle Net (SQL*Net) access. You will need the following information from the *Welcome to Oracle Cloud* email you received when the Cloud service with database schema information was activated:

- Service Home: the Cloud service URL
- User: the SFTP user for this Cloud service
- Secure FTP Site: the SFTP site URL

When you create a Cloud connection, you can use it immediately within the SQL Developer graphical interface; however, to use that connection with the SQL Developer command line interface, you must first exit the SQL Developer session in which you created the connection.

Connection Name: A name for this connection to a specified Cloud Database Service instance. Can be any name that you choose.

Database

Username: Username required during sign in when launching the Cloud service. (You cannot specify the password here. When you attempt to connect to the Cloud connection, an **Authentication** dialog box is displayed, prompting you for the user name and password.)

URL: Service Home URL from the *Welcome to Oracle Cloud* email under Service Details. May be in the following form:

`https://database-<identity-group>.db.<data-center>.<host-name>/<server-path>/`

Advanced: Displays the [Advanced Connection Information](#) (page 6-18) dialog box.

SFTP

Username: Secure FTP user name that you received in the *Welcome to Oracle Cloud* email under Service Details. (You cannot specify the password here. You will be prompted for the password when necessary.)

Hostname: Secure FTP Site from the *Welcome to Oracle Cloud* email under Service Details.

Port: Secure FTP port number (for example, 22).

6.31.1 Advanced Connection Information

This dialog box is displayed if you click Advanced in the New/Edit Cloud Connection dialog box.

HTTP or HTTPS: Select the desired protocol. (HTTPS provides secure communication.)

Hostname: *<host-name>* part of the URL in the New/Edit Cloud Connection dialog box.

Port: Port number (if any).

Server Path: *<server-path>* part of the URL in the New/Edit Cloud Connection dialog box. For example: /apex/

Service Name: Service name. For example: _sqldev/

Related Topics:

[New/Edit Cloud Connection](#) (page 6-18)

6.32 New/Edit NoSQL Connection

Use this dialog box to create or edit a connection to an Oracle NoSQL database.

Connection Name: A name for this connection to a Oracle NoSQL database. Can be any name that you choose.

Store

Username: User name on the NoSQL store.

Password: Password for the specified user.

Host: Host name for the NoSQL store.

Port: Port number on the specified host.

Store: Name of the NoSQL store. (*kvstore* is a typical KVLite store name.)

Security

SSL: Specifies the use of Secure Sockets Layer (SSL) encryption and network port restrictions, which enhance protection from network intrusion.

TrustStore: Location of the Java *truststore* file that is referenced by the *SSL_TRUSTSTORE_FILE_PROPERTY* property. You can click **Browse** to select the TrustStore file.

Oracle Wallet: Specifies Oracle Wallet as the password storage mechanism. This option is only available in the Oracle NoSQL Database EE version. You can click **Browse** to select the Wallet file.

Properties File: The security file that contains property settings for the login. You can click **Browse** to select the properties file.

6.33 Rename Model (Migration)

This dialog box is displayed when you right-click a captured or converted model and select Rename Model. To rename the model, change the name and click **OK**.

6.34 Delete Confirmation

This dialog box is displayed in certain situations to confirm whether you want to delete the selected object or objects. To perform the deletion, click **Yes**; to cancel the deletion request, click **No**.

6.35 Delete Confirmation (Migration)

This dialog box is displayed when you right-click a migration repository object under a captured or converted model and select Delete. To delete the object, click **Yes**.

If you do not want to be asked to confirm deletions of migration repository objects in the future, enable (check) **Skip This Message Next Time**. This will cause future deletions to occur when you right-click and select Delete.

6.36 Rename Database Item (Migration)

This dialog box is displayed when you right-click a database object under a captured or converted model and select Rename. To rename the object, change the name and click **OK**.

6.37 Select Connection

Use this dialog box to select a database or other type of connection for use with a specific SQL Developer feature (for example, the SQL worksheet, the Data Miner navigator, or the Reports navigator). After you click **OK**, the interface for the component is displayed, with the current user the same as the one specified in the connection.

To create a new connection, click the plus (+) icon; to edit the selected connection, click the pencil icon. In both cases, a dialog box for specifying connection information is displayed.

Related Topics:

[Create/Edit/Select Database Connection](#) (page 6-11)

6.38 Connection Information

Use this dialog box to specify the user name and password for the selected database connection.

If the specified user name does not exist in the database associated with the connection, or if the specified password is not the correct one for that user, the connection is refused.

Related Topics

- [Database Connections](#) (page 1-28)

6.39 No Connection Found

This dialog box is displayed when you attempt to perform an operation that requires a database connection, but no connection currently exists for that operation. For example, you might have opened a SQL file but not selected a connection, or the connection might have disconnected; or you might have tried to perform a schema copy operation without specifying both the From Schema and To Schema connections.

To select a connection in the SQL Worksheet, click **OK** to close this dialog box, then select a connection from the drop-down list in the SQL Worksheet icon bar.

6.40 Connection Rename Error

This dialog box is displayed when you attempt to rename a database connection to a name that is already used for another connection. For example, you might have forgotten to enter a new name for the connection that you want to rename.

To rename the connection, click OK to close this dialog box, then specify a unique connection name.

6.41 New Folder (Connections)

This dialog box enable you to create or rename a folder for organizing database connections. If you are creating a folder, enter the name of the new folder; if you are renaming a folder, replace the existing name with the desired new name. For information about using folders, see [Using Folders to Group Connections](#) (page 1-31).

6.42 Continue After Pause

This dialog box is displayed when a PAUSE statement is encountered in a script that you are running in the SQL Worksheet.

To continue execution at the statement after the PAUSE statement, click **Yes**. To stop execution and not continue with the statement after the PAUSE statement, click **No**.

6.43 Select Library

This dialog box is displayed when you click Browse in the Database pane when setting SQL Developer preferences. Use this box to select the library for the specified JDBC driver class.

Related Topics:

[Database](#) (page 1-99)

[SQL Developer Preferences](#) (page 1-92)

6.44 Create Library

This dialog box is displayed when you click New in the Select Library dialog box, which is displayed when you click Browse in the Database pane when setting SQL Developer preferences. Use this box to create the library for the specified JDBC driver class.

Related Topics:

[Database](#) (page 1-99)

[SQL Developer Preferences](#) (page 1-92)

6.45 CVS: Check Out from CVS

(Applies only if you have added support for CVS.)

Use this dialog box to check out modules from a CVS repository.

Connection Name: Name of the connection to the repository

Module Name: Name of the module to be checked out.

Path: Path to the module.

Get/Refresh Module List: Displays the list of modules or updates the current display.

Destination Folder: Folder into which to place the checked out files.

Use Revision or Tag: If this option is checked, the revision or tag that you specify in the text box is used. To see the available tags, click the binoculars icon.

Prune Empty Folders: If this option is checked, empty folders are removed from the working directory.

6.46 CVS: Create/Edit CVS Connection

(Applies only if you have added support for CVS.)

This information applies to creating or editing a CVS (Concurrent Versions System) connection.

Connection

Access Method: The method by which the client will gain access to and authenticate against the server. The methods available depend on which CVS preferences you have set; the available methods might include External, Password Server, Secure Shell via SSH2, and [Other].

Most of the remaining Connection fields apply only to specific Access Method values.

User Name: A CVS user name known to the repository.

Host Name: Qualified host name or IP address of the CVS server system.

Port: TCP/IP port number on which the repository is listening.

Repository Path: The location of the CVS repository software. The seeded / can be overwritten with a path in the format suitable for your operating system, for example c:\cvs. A simple formatting error, such as a forward slash instead of a backslash, will result in a message asking you to enter a valid repository path.

SSH2 Key: Path and file name for the SSH2 private key file for this connection. You can generate a SSH2 private key file using Generate SSH2 Key Pair.

Generate SSH2 Key Pair: Displays a dialog box for generating an SSH2 key pair (that is, a private key file and a public key file). You specify the private key file in the SSH2 Key box. You add the details of the public key file to the list of public keys on the CVS server system

Use HTTP Proxy Settings: Check (enable) this option if you are behind a firewall and need to use HTTP to access the CVS server.

External Locator Configuration: Displays the [External Locator Configuration](#) (page 6-87) dialog box, in which you can edit the details of the remote shell client and remote server program.

Root

Value of CVSROOT: CVS root variable made up from the information that you have already provided. This variable provides the client with access details when contacting the server. The format of the seeded variable

is: :accessmethod:username@serverlocation:repositorypath

You would not normally need to change this value. One instance when you would change this value is when you are attempting to connect to a CVSNT server through a firewall. In this case, you would add proxy information to the beginning of the username portion, so that the CVS root variable would take the following form: :accessmethod:proxy=proxyname;proxyport=portnumber:username@serverlocation:repositorypath

Connection

Test Connection: Attempts to establish a connection to the CVS repository.

Status: Displays the result of the test (success or an error message).

Name

Connection Name: Name to identify the connection to the CVS repository. The default name is the same as the CVSROOT value.

Summary

Displays the connection information that you have specified. To make any corrections, click **Back** as needed and modify the information. To create the connection, click **Finish**.

6.47 CVS: Import to CVS

(Applies only if you have added support for CVS.)

This interface is displayed when you click Versioning, then CVS, then Import Module. It enables you to import local files into the repository as a CVS module.

Module

Select the connection name, enter a name for the module, and optionally enter a descriptive comment about the import operation.

Tags

Select the connection name, enter a name for the module, and optionally enter a descriptive comment about the import operation.

Sources

Source Folder: Location from which files will be copied for the import operation.

Filters

You can configure filters to be used for excluding folders and files from the import operation. Use the arrow keys to move selected filters or all filters between Available Filters and Selected Filters.

To create a filter and add it to the Selected Filters list, click **New** to display the [Create Filter](#) (page 6-31) dialog box.

Options

You can specify options to be used during the import operation.

Use File Modification Time: If this option is checked, the file's modification time is used as the time of import. If this option is not checked, the time when the import operation is performed is used as the time of import.

Perform Module Checkout: If this option is checked, the modules are checked out after they are imported.

Summary

You can review the information that will be used to perform the import operation.

To go back and make any changes, click **Back** as needed.

To perform the import operation, click **Finish**.

6.48 CVS: Log In to CVS

(Applies only if you have added support for CVS.)

Use this dialog box to log in to the specified CVS repository. You must know the password for the specified user.

Connect Automatically on Startup: If this option is checked, a login operation is performed when you start SQL Developer.

6.49 Data Import Wizard (Load Data)

Use this wizard to import data into a table. For example, if you right-click the Tables node or a table name in the Connections navigator and select Import Data, you can specify the .source file (such as a spreadsheet or a delimited file) from which to import data. You create a table and import data into it, or import data into an existing table.

6.49.1 Data Preview

You can specify preferences that affect the preview display of data to be imported. Several default values are determined by the [Database](#) (page 1-99): Utilities: Import user preferences.

Header: If this option is enabled, a header row (not data to be imported) is assumed to start before (Before Skip) or after (After Skip) the number of rows for Skip Rows.

Skip Rows: The number of rows at the start to skip (that is, not be considered as data to be imported). Thus, the combination of Header and Skip Rows determines the total number of rows at the start that are not considered data to be imported.

Format: Format of the file containing data to be imported. For example: .xls (Microsoft Excel), .csv (comma separated value), .dsv (delimiter separated value), or .tsv (tab separated value).

Preview Limit: The maximum number of rows of data to be displayed in the preview pane in the lower part of the box. When the wizard creates a new table, the preview data is used to calculate the size of the columns; therefore, ensure that the preview is a good sample of the data. The displayed data is affected by the Preview Limit and by the [Database](#) (page 1-99): Utilities: Import user preference for Preview File Read Maximum, which limits the total number of bytes read from the file.

Encoding: Character set used for encoding of the data to be imported.

Delimiter, Line Terminator, Left Enclosure, Right Enclosure: Select or type the character in the input file that is used for each of these.

6.49.2 Import Method

Import Method: One of the following methods for loading the table definitions and data:

- **Insert:** For new tables, the table will be created and data will be inserted.
- **Insert Script:** A script will be generated with DDL statements to create the new tables and INSERT statements to add the data rows.
- **External Table:** External table DDL statements will be generated for accessing the file as a read-only table.

- **Staging External Table:** A script will be generated with DDL statements to create new tables and a staging external table for reading the file, and with INSERT statements for inserting the data into the table from the staging external table.
- **SQL*Loader:** For new tables, files will be created for running SQL*Loader to load the data; each table can be created during the import, or a script can be generated with DDL statements to create the table.

The number of remaining fields on this page and their availability depend on the load method and whether or not you invoked the wizard on a specific table.

Table Name: Target table in which to import the data.

Staging Table Name: External table to be used for loading the target table.

Commit and Drop Staging Table: If this option is enabled, the staging table is committed and dropped after the import operation. If this option is not enabled, the staging table is not committed or dropped.

Send Create Script to Worksheet: If this option is enabled, after you click Finish SQL Developer does not immediately perform the import operation, but instead opens a SQL Worksheet with statements that will be used if you click the Run Script icon in the worksheet.

6.49.3 Choose Columns

Available Columns: Lists the columns from which you can select for import into columns in the table. To select one or more columns, use the arrow buttons to move columns from Available to Selected.

Selected Columns: Lists the columns whose data is to be imported into columns in the database table. To change the order of a selected column in the list for the import operation, select it and use the up and down arrow buttons.

6.49.4 Column Definition

Enables you to specify information about the columns in database table into which to import the data.

Match By: The kind of automatic matching from source to target columns: by **Name**, by **Position**, or **None** (use the default). The default for existing tables is by name if a header is present in the file, and by position if no header is present.

Source Data Columns and **Target Table Columns:** You can select a source file data column to display its target (Oracle) column properties. For **Data Type**, select one of the supported types for this import operation. For a VARCHAR2 or NUMBER column, you must specify an appropriate **Size/Precision** value. You can specify whether the column value can be null (**Nullable?**), and you can specify a default value (**Default**).

6.49.5 Finish

Verify Parameters Before Import: You are encouraged to verify the import parameters. If any test fails, the Information column contains a brief explanation, and you can go back and fix any errors before you can click Finish.

To perform the import operation, click **Finish**. The action taken will depend on the load method selected and on whether you specified to send the script statements to a SQL Worksheet.

Related Topics

- [Database Objects](#) (page 1-13)

6.50 Export/Import Connections

The Export Connections wizard exports information about one or more database connections to an XML file. The Import Connections wizard imports connections that have been exported. Database connections that you import are added to any connections that already exist.

- [Export Connections](#) (page 6-26)
- [Import Connections](#) (page 6-26)

6.50.1 Export Connections

Select Connections

You can select and deselect all connections or specific connections for the operation.

Destination File

File Name: Name of the XML file to contain definitions of the connections to be exported. Use the Browse button to specify the location.

Password Handling

Specify how passwords should be handled in the exported file:

- **Encrypt all passwords with a key:** Specify an Encryption Key value and verify that value. (Anyone who attempts to import connections from the exported file will need to know the encryption key in order to be able to use connections with saved passwords without being prompted for the password.)
- **Remove all passwords from the exported connections:** Removes any saved passwords from the exported connections. (If the connections are later imported, users will need to know the passwords for connections that they plan to use.)

Summary

Displays a summary of the options you specified. To make any changes, press **Back** as needed and change the information. To start the operation, click **Finish**.

6.50.2 Import Connections

Source File

File Name: Name of the XML file that contains definitions of the connections to be imported. Use the Browse button to specify the location.

Password Handling

Either specify the key that was used to encrypt the passwords when the connections were exported, or remove all passwords for the imported connections:

- **Use a key to decrypt all passwords:** Specify the **Encryption Key** value that was used to encrypt the passwords for the export operation. (If you specify an incorrect encryption key, this option has the same effect as the option to remove all passwords during the import operation.)
- **Remove all passwords from the exported connections:** During the import operation, removes any passwords that are saved in exported connections. (Users

of the imported connections will need to know the passwords for connections that they plan to use.)

Select Connections

You can select and deselect all connections or specific connections for the operation.

Duplicate Connections: Determines what happens each existing connection that has the same name as a connection in the source file used for the import operation.

- **Rename:** Gives each duplicate connection a new name similar to the name of the existing connection.
- **Replace:** Replaces the information for the existing connection with the information for that imported connection.

Summary

Displays a summary of the options you specified. To make any changes, press **Back** as needed and change the information. To start the operation, click **Finish**.

Related Topics

- [Database Connections](#) (page 1-28)

6.51 Create/Edit Database Destination

This interface is used for creating or editing an Oracle Scheduler database destination.

(To create a database destination, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_DATABASE_DESTINATION procedure.)

Name: Name of the database destination. It can optionally be prefixed with a schema name. Cannot be null. It is converted to uppercase unless enclosed in double-quotes.

Enabled: Enables the database destination. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the database destination is created.)

Description: Optional text string that can be used to describe the database destination.

Connection Info: An Oracle Net connect identifier that is resolved to the Oracle database instance being connected to.

Agent: The external destination name of the Scheduler agent to connect. Equivalent to an agent name.

Related Topics

- [Database Destinations](#) (page 1-15)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.52 Create/Edit Destination Group (Database or External)

This interface is used for creating a new Oracle Scheduler database or external destination group.

Properties

Name: Name of the destination group. It can optionally be prefixed with a schema name. Cannot be null. It is converted to uppercase unless enclosed in double-quotes.

Enabled: Enables the destination group. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the destination group is created.)

Description: Optional text string that can be used to describe the destination group.

Members

Select members for the group, and specify whether to use the default credential for all members.

Related Topics

- [Destination Groups](#) (page 1-15)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.53 Create/Edit Database Link

The following information applies to a database link, which is a database object in one database that enables you to access objects on another database, as explained in [Database Links \(Public and Private\)](#) (page 1-15).

Public: If this option is checked, the database link is public (available to all users). If this option is not checked, the database link is private and is available only to you.

Schema: Database schema in which to create the database link.

Name: Name of the database link. Must be unique within a schema.

Host Name: The service name of a remote database. If you specify only the database name, Oracle Database implicitly appends the database domain to the connect string to create a complete service name. Therefore, if the database domain of the remote database is different from that of the current database, you must specify the complete service name.

Current User: Creates a current user database link. The current user must be a global user with a valid account on the remote database. If the database link is used directly, that is, not from within a stored object, then the current user is the same as the connected user.

Fixed User: Creates a fixed user database link, for which you specify the user name and password used to connect to the remote database.

Shared: If this option is checked, a single network connection is used to create a public database link that can be shared among multiple users. In this case, you must also specify the Authentication information.

Authentication - User Name and Password: The user name and password on the target instance. This information authenticates the user to the remote server and is required for security. The specified user and password must be a valid user and password on the remote instance.

DDL tab

You can review and save the SQL statement that SQL Developer will use to create the database link.

Related Topics

- [Database Objects](#) (page 1-13)

6.54 Create/Edit Index

The following information applies to an index, which is a database object that contains an entry for each value that appears in the indexed column or columns of the table or

cluster and provides direct, fast access to rows, as explained in [Indexes](#) (page 1-17). For detailed information about all index-related options, see the CREATE INDEX reference section in *Oracle Database SQL Language Reference*.

Schema: Database in which to create the index (often the same as the user creating the index).

Name: Name of the index. Must be unique within a schema.

6.54.1 Definition

Table Schema: Database schema that owns the table associated with the index.

Table: Name of the table associated with the index.

Index Type: The type of Oracle index. **Non-unique** means that the index can contain multiple identical values; **Unique** means that no duplicate values are permitted; **Bitmap** stores rowids associated with a key value as a bitmap; **Domain** lets you define an index for a specialized domain (for example, for a spatial index: Indextype Schema = MDSYS, Indextype = SPATIAL_INDEX).

Expressions: A list of index expressions, that is, the table columns or column expressions in the index. To add an index expression, click the Add Column Expression (+) icon; this adds a column name here and in Column Expression, where you can edit it. To delete an index expression, click the Remove Column Expression (X) icon; to move an index expression up or down in the list, click the Move Column Up and Move Column Down icons. An index must have at least one index expression.

For example, to create an index on the AUTHOR_LAST_NAME column of the BOOKS table from the tutorial, click the + icon, and select AUTHOR_LAST_NAME in Column Name or Expression (next field), which changes BOOKS to AUTHOR_LAST_NAME in the Index field.

Expression: A column name or column expression. A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.

Order: ASC for an ascending index (index values sorted in ascending order); DESC for a descending index (index values sorted in descending order).

Related Topics:

[Create a Table \(BOOKS\)](#) (page 5-2)

6.54.2 Properties

Enables you to specify index properties

Key Compression: You can enable key compression is enabled, which eliminates repeated occurrence of key column values and may substantially reduce storage. If this option is checked, you can enter an integer to specify the prefix length (number of prefix columns to compress).

Parallel Degree and Degree: If you enable parallel creation of the index, specify `Default` for the default degree (the optimum degree of parallelism is automatically calculated) or `Select` to specify an integer value for the degree of parallelism, which is the number of parallel threads used in the parallel operation. (Each parallel thread may use one or two parallel execution servers.)

6.54.3 Storage

Storage Options: Lets you specify [Storage Options](#) (page 6-54) for the index.

6.54.4 Partitions

When applicable, enables you to specify whether the index is not partitioned, locally partitioned, or globally partitioned. If you specify Local or Global, additional fields are displayed relevant to the selected partitioning option.

Local: Specifies that the index is partitioned on the same columns, with the same number of partitions and the same partition bounds as its associated table. Oracle Database automatically maintains local index partitioning as the underlying table is repartitioned.

Global: Specifies that the partitioning of the index is user defined and is not equipartitioned with the underlying table. You can partition a global index by range or by hash. In both cases, you can specify up to 32 columns as partitioning key columns. The partitioning column list must specify a left prefix of the index column list. If the index is defined on columns a, b, and c, then for the columns you can specify (a, b, c), or (a, b), or (a, c), but you cannot specify (b, c) or (c) or (b, a). If you omit the partition names, then Oracle Database assigns names of the form SYS_P*n*.

6.54.5 DDL

Read-only display of the DDL statement or statements to create or edit the index using the current specifications.

Save: Click to save the DDL statement or statements to a .sql file.

Related Topics

- [Database Objects](#) (page 1-13)

6.55 Create File Watcher

This interface is used for creating a new Oracle Scheduler file watcher.

(To create a file watcher, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_FILE_WATCHER procedure.)

Name: Name of the file watcher. It can optionally be prefixed with a schema name. Cannot be null. It is converted to uppercase unless enclosed in double-quotes.

Enabled: Enables the file watcher. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the file watcher is created.)

Description: Optional descriptive text.

Destination: Name of an external destination. You create an external destination by registering a remote Scheduler agent with the database. The view ALL_SCHEDULER_EXTERNAL_DESTS lists valid external destination names. If this parameter is null, the file watcher is created on the local host.

Directory Path: Directory in which the file is expected to arrive. The single wildcard '?' at the beginning of the path denotes the Oracle home path. For example, '?/rdbms/log' denotes the rdbms/log subdirectory of the Oracle home directory.

File Name: Name of the file to look for. Two wildcards are permitted anywhere in the file name: '?' denotes any single character, and '*' denotes zero or more characters.

Credential Name: Name of a valid credential object. The file watcher uses the credential to authenticate itself with the host operating system to access the watched-for file. The file watcher owner must have EXECUTE privileges on the credential.

Min File Size: Minimum size in bytes that the file must be before the file watcher considers the file found.

Steady State Duration: Minimum time interval that the file must remain unchanged before the file watcher considers the file found. Cannot exceed one hour. If null, an internal value is used. The minimum value is 10 seconds. Oracle recommends similar steady state duration values for all file watchers for efficient file watcher job operation. Also, the repeat interval of the file watcher schedule must be equal or greater than the steady state duration value.

Related Topics

- [File Watchers](#) (page 1-16)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.56 Create Filter

This dialog box is displayed when you click New to add a user-defined exclusion filter when importing files into a repository.

Filter: Shell filename pattern, which can contain both normal characters and meta-characters, including wildcards. (See the supplied Selected Filters list for typical patterns.) For example, to exclude files with the extension *xyz*, enter the following:
*.xyz

When you click OK, the specified filter is added to the Selected Filters list.

6.57 Create Function or Procedure

Use this dialog box to create a PL/SQL subprogram (function or procedure).

Schema: Database schema in which to create the subprogram.

Name: Name of the subprogram. Must be unique within a schema.

Add New Source in Lowercase: If this option is checked, new text is entered in lowercase regardless of the case in which you type it. This option affects only the appearance of the code, because PL/SQL is not case-sensitive in its execution.

Parameters: Specify information about parameters. To add a parameter, click the Add Parameter (+ or sign) icon; to delete a parameter, select it and click the Remove Parameter (X) icon.

- **Name:** Name of the parameter.
- **Mode:** IN for input, OUT for output, or IN OUT for both input and output.
- **No Copy:** Requests that the compiler pass the corresponding actual parameter by reference instead of value. Each time the subprogram is invoked, the optimizer decides, silently, whether to obey or disregard NOCOPY.
- **Data Type:** Data type for the parameter, If the desired value is not included in the drop-down list, you can type it in the box.
- **Default Value:** Value used if no value is specified at run time.

After you click OK, the subprogram is created and is displayed in the Editor window, where you can enter the details.

Related Topics

- [Database Objects](#) (page 1-13)
- [Functions](#) (page 1-16)
- [Procedures](#) (page 1-20)
- [Packages](#) (page 1-20)

6.58 Create/Edit Job

This interface is used for creating a new Oracle Scheduler job or editing an existing job. For more information about job scheduling, see [Scheduling Jobs Using SQL Developer](#) (page 1-67).

(To create a job, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_JOB procedure, which is documented in *Oracle Database PL/SQL Packages and Types Reference*.)

- [Job Details](#) (page 6-32)
- [Destination](#) (page 6-33)
- [Job Arguments](#) (page 6-33)
- [Notification](#) (page 6-33)
- [Properties](#) (page 6-33)
- [Summary/SQL](#) (page 6-34)

6.58.1 Job Details

Job Name: Name of the job.

Enabled: If this option is specified, validity checks will be made and the job will be created enabled if all the checks are successful. If this option is not specified, the job is not created enabled.

Description: Optional text string that can be used to describe the job.

Job Class: Name of the job class to which this job belongs.

Type of Job: Type of object to be executed by the job: PL/SQL Block, Chain, Stored Procedure, Named Program, or Executable.

When to Execute Job: When to execute the job: Immediate (immediately on creation, and once only), Once (once, at a specified time), Repeating, Queue, File Watcher, Schedule (using a named schedule object). If you specify anything other than Immediate, you are prompted for additional information.

For Repeating, you can click the pencil icon to specify a **Repeat Interval** with details. If you check **Advanced**, you can use the predefined interval scheme or a user-defined one, or a combination.

6.58.2 Destination

Local (local system), **Remote** (the database destination for a remote database job, or external destination for a remote external job), or **Multiple** (the job runs on all destinations).

Depending on what destination you selected for the job, select the local credential, the remote credential and destination, or the multiple group.

6.58.3 Job Arguments

If the job expects arguments to be passed, the Name and Data Type of each argument are listed. For each input or input/output argument, specify the appropriate Value.

6.58.4 Notification

Select Job Events for which to send notification email messages. For each message, you can specify recipient email addresses and the sender (or no sender), and you can modify the body of the message.

6.58.5 Properties

Auto Drop: Determines whether the job is to be automatically dropped after it has completed or has been automatically disabled.

Restartable: Determines whether the job can be restarted in case of failure.

Stop on Window Close: If the schedule of a job is a window or a window group, TRUE causes the job to stop once the associated window is closed, and FALSE causes the job to continue after the window closes. (Note that if the job is allowed to continue, its resource allocation will probably change because closing a window generally also implies a change in resource plans.)

Follow Default Time Zone: Determines whether if the job start date is null, then when the default time zone scheduler attribute is changed, the Scheduler recomputes the next run date and time for this job so that it is in accordance with the new time zone.

Parallel Instances: For an event-based job, determined what happens if an event is raised and the event-based job that processes that event is already running. FALSE causes the new event to be ignored. TRUE causes an instance of the job to be started for every instance of the event, and each job instance is a lightweight job so multiple instances of the same event-based job can run in parallel.

Job Style: Style of the job being created: REGULAR (regular job) or LIGHTWEIGHT (lightweight job). A lightweight must reference a program object. Use lightweight jobs when you have many short-duration jobs that run frequently. Under certain circumstances, using lightweight jobs can deliver a small performance gain.

Job Priority: The priority of this job relative to other jobs in the same class as this job. If multiple jobs within a class are scheduled to be executed at the same time, the job priority determines the order in which jobs from that class are picked up for execution by the job coordinator. It can be a value from 1 through 5, with 1 being the first to be picked up for job execution.

Job Weight: (Do not change the value. Shown only for connections to Oracle Database releases before 11.2.)

Logging Level: Determines how much information is logged:
DBMS_SCHEDULER.LOGGING_OFF (no logging),

DBMS_SCHEDULER.LOGGING_FAILED_RUNS (only jobs that failed, with the reason for failure), DBMS_SCHEDULER.LOGGING_RUNS (all runs of each job in this class), or DBMS_SCHEDULER.LOGGING_FULL (all operations performed on all jobs).

However, if the job class has a higher (more detailed) logging level than the level specified for the job, the job class logging level is used.

Max Runs: The maximum number of consecutive scheduled runs of the job.

Max Failures: The number of times a job can fail on consecutive scheduled runs before it is automatically disabled.

Instance ID: In an Oracle Real Application Clusters environment, the instance ID of the instance that the job must run on.

Raise Events: Determines at what stages of the job execution to raise events.

Max Run Duration: Maximum amount of time that the job should be allowed to run. Its data type is INTERVAL DAY TO SECOND. If this attribute is set to a nonzero and non-null value, and job duration exceeds this value, the Scheduler raises an event of type JOB_OVER_MAX_DUR. It is then up to your event handler to decide whether or not to allow the job to continue.

Schedule Limit: Maximum delay time between scheduled and actual job start before a program run is canceled.

Store Output: If set to TRUE, then for job runs that are logged, all job output and error messages are stored in the *_JOB_RUN_DETAILS views. If set to FALSE, then the output and messages are not stored.

Reset to Defaults: Resets all properties to their default values.

6.58.6 Summary/SQL

May include a summary in hierarchical form or the PL/SQL statement that will be used to implement your specifications, or both.

Related Topics

- [Jobs](#) (page 1-18)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.59 Create/Edit Job Class

This interface is used for creating a new Oracle Scheduler job class or editing an existing job class.

(To create a job class, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_JOB_CLASS procedure, which is documented in *Oracle Database PL/SQL Packages and Types Reference*.)

Name: Name of the job class. If you specify a schema, it must be SYS. For an existing job class, this field is read-only; to change the name, you must drop the job class and create a new job class with the desired name.

Description: Optional text string that can be used to describe the job class.

Logging Level: Specifies how much information is written to the job log: **RUNS** (detailed information for all runs of each job in this class), **FULL** (RUNS plus all operations performed on all jobs in this class), or **OFF** (no logging).

Log Retention Period (days): Number of days that job log entries for jobs in this class are retained. The range of valid values is 0 through 999. If set to 0, no history is kept. If NULL (the default), retention days are set by the `log_history` Scheduler attribute (set with `SET_SCHEDULER_ATTRIBUTE`).

Service Name: The database service that the jobs in this class will have affinity to. If no service is specified, the job class will belong to the default service, which means it will have no service affinity and any one of the database instances within the cluster might run the job.

Resource Consumer Group: Resource consumer group this class is associated with. If no resource consumer group is specified, the job class is associated with the default resource consumer group.

Related Topics

- [Job Classes](#) (page 1-18)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.60 Create/Edit Materialized View Log

User this dialog box to create or edit a materialized view log, which is a table associated with the master table of a materialized view. For more information, see [Materialized View Logs](#) (page 1-19).

Schema: Database schema in which to create the materialized view log.

Name: Name of the master table of the materialized view to be associated with this materialized view log.

Properties tab

Tablespace: Tablespace in which the materialized view log is to be created.

Logging: LOGGING or NOLOGGING, to establish the logging characteristics for the materialized view log.

Row ID: Yes indicates that the rowid of all rows changed should be recorded in the materialized view log; No indicates that the rowid of all rows changed should not be recorded in the materialized view log.

Primary Key: Yes indicates that the primary key of all rows changed should be recorded in the materialized view log; No indicates that the primary key of all rows changed should not be recorded in the materialized view log.

New Values: INCLUDING saves both old and new values for update DML operations in the materialized view log; EXCLUDING disables the recording of new values in the materialized view log. If this log is for a table on which you have a single-table materialized aggregate view, and if you want the materialized view to be eligible for fast refresh, you must specify INCLUDING.

Cache: For data that will be accessed frequently, CACHE specifies that the blocks retrieved for this log are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. NOCACHE specifies that the blocks are placed at the least recently used end of the LRU list.

Parallel: If this option is checked, parallel operations will be supported for the materialized view log.

Object ID: For a log on an object table only: Yes indicates that the system-generated or user-defined object identifier of every modified row should be recorded in the materialized view log; No indicates that the system-generated or user-defined object identifier of every modified row should not be recorded in the materialized view log.

Sequence: Yes indicates that a sequence value providing additional ordering information should be recorded in the materialized view log; No indicates that a sequence value providing additional ordering information should not be recorded in the materialized view log. Sequence numbers (that is, Yes for this option) are necessary to support fast refresh after some update scenarios.

Available Filter Columns: Additional columns, which are non-primary-key columns referenced by subquery materialized views, to be recorded in the materialized view log. To select one or more filter columns, use the arrow buttons to move columns from Available to Selected.

DDL tab

You can view a SQL CREATE statement that reflects the current definition of the object, or a SQL ALTER statement to modify an existing object to reflect your changes.

To save the SQL statement to a script file, click **Save** and specify the location and file name.

Related Topics

- [Materialized View Logs](#) (page 1-19)
- [Create/Edit View](#) (page 6-62) (information related to materialized views)

6.61 Create PL/SQL Package

Use this dialog box to create a package to contain PL/SQL subprograms (functions or procedures, or a combination).

Schema: Database schema in which to create the PL/SQL package.

Name: Name of the package. Must be unique within a schema.

Add New Source in Lowercase: If this option is checked, new text is entered in lowercase regardless of the case in which you type it. This option affects only the appearance of the code, because PL/SQL is not case-sensitive in its execution.

The package is created and is displayed in the Editor window, where you can enter the details.

Related Topics

- [Packages](#) (page 1-20)

6.62 Create Pluggable Database

Use this dialog box to create a PDB.

Database Name: Name of the PDB.

Admin Name: Administrative user who has or can be granted the privileges required to perform administrative tasks on the PDB.

Admin Password: Password for the Admin Name database user.

File Name Conversions: Determines how the database generates the names of files (such as data files and wallet files) for the PDB.

- **None:** The database first attempts to use Oracle Managed Files to generate file names. If you are not using Oracle Managed Files, then the database uses the PDB_FILE_NAME_CONVERT initialization parameter to generate file names.
- **Custom Names:** Select a Source Files/Target Files pair.
- **Custom Expressions:** Specify one or more Source File Expression/Target File Expression pairs. Each pair item is a string found in names of files associated with the seed (when creating a PDB by using the seed), associated with the source PDB (when cloning a PDB), or listed in the XML file (when plugging a PDB into a CDB).

Storage: You can specify storage limits for the PDB total size or temporary tablespace usage, or both; or unlimited storage for either or both.

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.63 Create/Edit Program

This dialog box is used for creating or editing an Oracle Scheduler program, which can include creating a new program using details from an existing program (Create Like). For more information about job scheduling, see [Scheduling Jobs Using SQL Developer](#) (page 1-67).

(To create a program, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_PROGRAM procedure.)

Name: Name of the program. The name has to be unique in the SQL namespace. For example, a program cannot have the same name as a table in a schema.

Enabled: If this option is specified, validity checks will be made and the program will be created enabled if all the checks are successful. If this option is not specified, the program is not created enabled.

Description: Optional text string that can be used to describe the program.

Type of program:

- **PL/SQL Block:** The program is a PL/SQL block. Job or program arguments are not supported when the job or program type is PLSQL_BLOCK. In this case, the number of arguments must be 0. Enter or paste in the complete PL/SQL code, or edit the existing code.
- **Stored Procedure:** The program is a PL/SQL or Java stored procedure, or an external C subprogram. Only procedures, not functions with return values, are supported. PL/SQL procedures with INOUT or OUT arguments are not supported.

Schema: Schema of the stored procedure. If not specified, the schema of the job is assumed.

Procedure: Name of the stored procedure.

Arguments: For each argument: name, data type, default value, and whether it is an input, output, or input/output argument.

- **Remote Stored Procedure:** The program is a remote stored procedure. You specify the procedure name and any arguments.

Procedure: Name of the stored procedure.

Arguments: For each argument: name, data type, default value, and whether it is an input, output, or input/output argument.

- **Executable:** The program is external to the database. External programs implies anything that can be executed from the operating system's command line. AnyData arguments are not supported.

Executable: Name of the external executable, including the full path name, but excluding any command-line arguments. If the action starts with a single question mark (?), the question mark is replaced by the path to the Oracle home directory for a local job or to the Scheduler agent home for a remote job. If the action contains an at-sign (@) and the job is local, the at-sign is replaced with the SID of the current Oracle instance.

Arguments: For each argument: name, data type, default value, and whether it is an input, output, or input/output argument

- **Script:** The program is a **SQL Script** (SQL*Plus statements), **Backup Script** (RMAN commands), or **External Script** (operating system commands). Enter or paste the script text in the box.

Properties tab

Enables you to set program properties. For most properties the default is null, but you can check the box to specify a value.

Detached: TRUE if the program is a detached job; otherwise, FALSE. Use a detached job to start a script or application that runs in a separate process, independently and asynchronously to the Scheduler. A detached job typically starts another process and then exits. Upon exit (when the job action is completed) a detached job remains in the running state. The running state indicates that the asynchronous process that the job started is still active. When the asynchronous process finishes its work, it must connect to the database and call DBMS_SCHEDULER.END_DETACHED_JOB_RUN, which ends the job.

Max Runs: Maximum number of runs before the program is marked as completed.

Max Failures: Maximum number of failures tolerated before the program is marked as broken.

Max Run Duration: Maximum run duration of the program.

Schedule Limit: Maximum delay time between scheduled and actual job start before a program run is canceled.

Related Topics

- [Programs](#) (page 1-21)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.64 Create/Edit Role

This dialog box is used for creating a new database role or editing an existing database role. The role properties are grouped under several tabs.

To create or edit a role, the user associated with your database connection must have the DBA role. You should also be familiar with the main concepts and techniques documented in *Oracle Database Administrator's Guide* and in the chapter about configuring privilege and role authorization in *Oracle Database Security Guide*.

Role tab

Specifies general properties for the role.

Role Name: The role name string. For an existing role, this field is read-only; to change the name, you must drop the user and create a new user with the desired name.

New Password: If you specify a password, the role is a *local role* and a user must specify the password to the database when enabling the role. If you do not specify a password, the role is authorized by the database and no password is required to enable the role.

If you specify a password, you must type the same password string for **Confirm Password**.

Granted Roles tab

Specifies roles to be granted to the role being created or modified. For each listed role, you can check **Granted** to grant the specified role to this role, **Admin** to permit the role to grant this role to other users or roles, and **Default** to use the default settings for Granted and Admin.

For convenience, you can click buttons to affect all settings (Grant All, Revoke All, Admin All, Admin None, Default All, Default None); then, you can specify other settings for individual roles.

System Privileges tab

Specifies privileges to be granted to the user. For each privilege, you can check **Granted** to grant the privilege, and **Admin Option** to permit the user to grant the privilege to other users.

For convenience, you can click buttons to affect all settings (Grant All, Revoke All, Admin All, Admin None); then, you can specify other settings for individual privileges.

SQL tab

Displays the SQL statements that SQL Developer will use to create (after executing a CREATE ROLE statement) a new user or to edit an existing role. This display is read-only; if you want to make any changes, go back to the relevant tabs and make the changes there.

Related Topics

- [Using DBA Features in SQL Developer](#) (page 1-54)

6.65 Create/Edit Schedule

This dialog box is used for creating or editing an Oracle Scheduler schedule.

(To create a schedule, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_SCHEDULE procedure.)

Properties tab

Name: Name of the schedule. The name has to be unique in the SQL namespace. For example, a schedule cannot have the same name as a table in a schema.

Description: Optional text string that can be used to describe the schedule.

When: Specify when jobs that use this schedule are to run: **Repeating** (specify the repeat interval, start date, and end date), **Queue** (specify the queue name, agent, condition, start date, and end date), or **File Watcher** (specify the file watcher object name, condition, start date, and end date).

SQL tab

Displays the DBMS_SCHEDULER.CREATE_SCHEDULE procedure that will be used to create the schedule according to your specifications.

Related Topics

- [Programs](#) (page 1-21)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.66 Create/Edit Sequence

The following information applies to a sequence, which is an object from which multiple users may generate unique integers. You can use sequences to automatically generate primary key values.

Schema: Database schema in which to create the sequence.

Name: Name of the sequence. Must be unique within a schema.

Increment: Interval between successive numbers in a sequence.

Start with: Starting value of the sequence.

Min value: Lowest possible value for the sequence. The default is 1 for an ascending sequence and $-(10^{26})$ for a descending sequence.

Max value: Highest possible value for the sequence. The default is 10^{27} for an ascending sequence and -1 for a descending sequence.

Cache and Cache size: Cache causes sequence values are preallocated in cache, which can improve application performance; Cache size indicates the number of sequence values preallocated in cache. No Cache causes sequence values not to be preallocated in cache.

Cycle: Indicates whether the sequence "wraps around" to reuse numbers after reaching its maximum value (for an ascending sequence) or its minimum value (for a descending sequence). If cycling of values is not enabled, the sequence cannot generate more values after reaching its maximum or minimum value.

Order: Indicates whether sequence numbers are generated in the order in which they are requested. If No Order is specified, sequence numbers are not guaranteed to be in the order in which they were requested.

DDL tab

You can review the SQL statement that SQL Developer will use to create a new sequence or that reflects any changes you have made to the sequence properties.

Related Topics

- [Sequences](#) (page 1-22)

6.67 Create SQL File

Use this dialog box to create a SQL script file and to open the file in a SQL Worksheet for editing.

File Name: Name and extension of the file to be created. The default and recommended extension is .sql.

Directory Name: Directory path for the file. To specify a directory, you can click Browse. The default directory is the [Location of User-Related Information](#) (page 1-122).

6.68 Create/Edit Synonym

The following information applies to a synonym, which is an alternative name for a table, view, sequence, procedure, stored function, package, materialized view, Java class database object, user-defined object type, or another synonym.

Public: If this option is checked, the synonym is accessible to all users. (However each user must have appropriate privileges on the underlying object in order to use the synonym.) If this option is not checked, the synonym is a private synonym, and is accessible only within its schema.

Schema: Database schema in which to create the synonym.

Name: Name of the synonym. A private synonym must be unique within its schema; a public synonym must be unique within the database.

For - Referenced Schema: Schema containing the object or name to which this synonym refers.

Object Based: Specify the object to which this synonym refers.

Name Based: Enter the name of the object to which this synonym refers.

DDL tab

You can review the SQL statement that SQL Developer will use to create a new synonym or that reflects any changes you have made to the synonym properties.

Related Topics

- [Synonyms \(Public and Private\)](#) (page 1-23)

6.69 Create Table (quick creation)

This dialog box (if you do not check the Advanced box) creates a new table quickly by specifying columns and some frequently used features. (If you need to add or change features after you create the table, you can edit the table by clicking the Modify icon while viewing the table or by right-clicking its name in the Connections navigator and selecting Properties, which displays the Create/Edit Table (with advanced options) dialog box.)

To create a new table, the only things you **must** do are specify the schema and the table name, add the necessary columns, and click OK. Although it is not required, you should also specify a primary key.

Advanced: If this option is checked, the dialog box changes to include an extended set of features for creating the table. For example, you must check this option if you want to create a partitioned table, an index-organized table, or an external table.

Schema: Database schema in which to create the table.

Name: Name of the table. Must be unique within a schema.

Table tab (quick creation)

Specifies properties for each column in the table.

Columns: Lists the columns currently in the table.

Note:

To add a column after the currently selected column, click **Add Column**; to delete a column, select it and click **Remove Column**.

Column Name: Name of the column. Must be unique within the table. Suggestion: For a new column, replace any default name, such as COLUMN1.

Type: Data type for the column. The drop-down list includes only selected frequently used data types. To specify any other type for the column, you must use the Columns panel of the Create/Edit Table (with advanced options) dialog box.

Size: For VARCHAR2 data, the maximum size of the column data; for NUMBER data, the maximum number of digits.

Not Null: If this option is checked, the column must contain data; you cannot specify no value or an explicit null value for this column when you insert a row. If this option is not checked, the column can contain either data or no data.

Primary Key: If this option is checked, the column is the primary key, or part of the primary key, for the table. The primary key is the column, or set of columns, that uniquely identifies each row in the table. A primary key column cannot be null.

If you want to have the primary key values automatically populated by a convenient method that uses a before-insert trigger and a sequence, then before you finish creating the table, you must check the Advanced box and use the Primary Key tab, starting with the Populate Primary Key Column field.

To add another column, click **Add Column**. When you are finished adding columns, either click **OK** or click the DDL tab to review the CREATE TABLE statement.

DDL tab (quick creation)

You can review and save the CREATE TABLE statement that SQL Developer will use to create a new table or that reflects any changes you have made to the table properties. If you want to make any changes, go back to the Table tab and make the changes there.

When you are finished, click **OK**.

Related Topics

- [Create/Edit Table \(with advanced options\)](#) (page 6-42)
- [Tables](#) (page 1-23)

6.70 Create/Edit Table (with advanced options)

The table dialog box is used for creating a new table or editing an existing table. The table properties are grouped in several panes.

To create a new table, the only things you **must** do are specify the schema and the table name, add the necessary columns, and click OK. Although it is not required, you should also specify a primary key using the [Primary Key Constraint](#) (page 6-44). For other table-related features, use the appropriate tabs; the order in which you visit tabs usually does not matter, although you might find it convenient to visit them in the sequence in this topic. If you are editing an existing table, you can visit the tabs in any order.

If you click OK before you are finished creating or editing the table, right-click the table name in the Connections navigator, select Edit, and continue creating or editing the table.

Schema: Database schema in which to create the table.

Name: Name of the table. Must be unique within a schema.

Table Type: The type of table:

- **Normal:** A regular database table. It can be partitioned .
- **External:** An external table.
- **Index Organized:** An index-organized table.
- **Temporary Table:** A temporary table, which is not stored permanently in the database. The temporary table definition persists in the same way as the definition of a regular table, but the table segment and any data in the temporary table persist only for the duration of either the transaction (**Transaction** option) or the session (**Session** option).

The following panes may also be available. (Some panes are available only for tables of specific types or with specific features.)

[Columns pane](#) (page 6-43)

[Constraints](#) (page 6-44)

[Primary Key Constraint](#) (page 6-44)

[Unique Constraints](#) (page 6-45)

[Foreign Keys Constraints](#) (page 6-46)

[Check Constraints](#) (page 6-46)

[Indexes pane](#) (page 6-47)

[In-Memory pane](#) (page 6-47)

[Storage pane](#) (page 6-48)

[Column Sequences pane](#) (page 6-48)

[Table Properties pane](#) (page 6-49)

[LOB Parameters pane](#) (page 6-49)

[Partitioning pane](#) (page 6-50)

[Subpartition Templates pane](#) (page 6-51)

[Partition Definitions pane](#) (page 6-51)

[Columns pane](#) (page 6-43)

[External Table Properties pane](#) (page 6-52)

[Comment pane](#) (page 6-53)

[DDL pane](#) (page 6-53)

6.70.1 Columns pane

Specifies properties for each column in the table.

Columns: Lists the columns currently in the table. To add a column, click the Add Column (+) icon; to delete a column, select it and click the Remove Column (X) icon; to move a column up or down in the table definition, select it and use the up-arrow and down-arrow buttons.

To copy in column definitions from another table, click the Copy Columns icon to display the [Copy Columns](#) (page 6-6) dialog box.

Note:

After you add a column, **to add another column**, click the Add Column (+) icon.

Name: Name of the column. Must be unique within the table. Suggestion: For a new column, replace any default name, such as COLUMN1.

Datatype: **Simple** indicates a simple (non-object) data type; **Complex** indicates an object type. For a complex type, you must specify the schema and the type name (for example, MDSYS and SDO_GEOMETRY for the Oracle Spatial and Graph geometry type).

Type: Name of the data type. Most of the remaining information depends on the specific type.

Precision: For numeric data, the precision (total number of significant digits that can be represented) of the column data.

Scale: For numeric data, the scale (number of digits after the decimal point) of the column data.

Size: For character data, the maximum size of the column data.

Units: For character data, the units represented by the Size: **BYTE** for bytes or **CHAR** for characters. This attribute is important if the database can contain data in Unicode format, with multiple bytes for each character.

Default: For relevant types, the default value inserted into the column if no value is specified when a row is inserted.

Cannot be NULL: If this option is checked, the column must contain data; you cannot specify no value or an explicit null value for this column when you insert a row. If this option is not checked, the column can contain either data or no data. A primary key column cannot be null.

Comment: Optional descriptive comment about the column.

To add another column, click the Add Column (+) icon.

6.70.2 Constraints

You can click the plus sign (+) icon to add one or more constraints. For each, specify the type of constraint: Primary Key, Unique, Foreign Key, or Check. (See the information for each type.)

6.70.3 Primary Key Constraint

Specifies the primary key for the table. The primary key is the column, or set of columns, that uniquely identifies each row in the table.

An index is automatically created on the primary key.

Name: Name of the constraint to be associated with the primary key definition. Must be unique within the database.

Enabled: If this option is checked, the primary key constraint is enforced: that is, the data in the primary key column (or set of columns) must be unique and not null.

Available Columns: Lists the columns that are available to be added to the primary key definition.

Selected Columns: Lists the columns that are included in the primary key definition.

To add a column to the primary key definition, select it in Available Columns and click the Add (>) icon; to remove a column from the primary key definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the primary key definition, select it in Selected Columns and use the arrow buttons.

The remaining fields (Populate Primary Key Column through Trigger Name) appear *only* when you are creating a table. They are not available when you are editing an existing table.

Populate Primary Key Column: When you are creating a table, if you want to use a trigger and a sequence to have a unique value automatically inserted into the primary key column when you insert a new row, specify the primary key column.

From: An existing sequence that you select, or a new sequence whose name you enter. (For a new sequence, SQL Developer creates the sequence automatically using the name that you enter.)

Trigger Name: The name for the before-insert trigger that will be automatically created. This trigger uses the sequence to generate a new value for the primary key when a row is inserted. For an example of using this technique, see the tutorial section [Create a Table \(TRANSACTIONS\)](#) (page 5-5).

6.70.4 Unique Constraints

Specifies one or more unique constraints for the table. A unique constraint specifies a column, or set of columns, whose data values must be unique: each data value must not be null, and it must not be the same as any other value in the column.

For a multicolumn unique constraint, the combination of values must be unique, and no column in the constraint definition can have a null value. For example, if you specify the office_name and city columns for a unique constraint, you could not have two Sales offices in Chicago, but you could have a Sales office in Chicago and a Sales office in Atlanta.

Unique Constraints: Lists the unique constraints currently defined on the table. To add a unique constraint, click the Add button; to delete a unique constraint, select it and click the Remove button.

Note:

After you add a unique constraint, **to add another unique constraint**, click the **Add** button.

Name: Name of the unique constraint. Must be unique within the database.

Enabled: If this option is checked, the unique constraint is enforced.

Available Columns: Lists the columns that are available to be added to the unique constraint definition.

Selected Columns: Lists the columns that are included in the unique constraint definition.

To add a column to the unique constraint definition, select it in Available Columns and click the Add (>) icon; to remove a column from the unique constraint definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the unique constraint definition, select it in Selected Columns and use the arrow buttons.

6.70.5 Foreign Keys Constraints

Specifies one or more foreign keys for the table. A foreign key specifies a column ("local column"), each of whose data values must match a value in the primary key or unique constraint of another table.

Foreign Keys: Lists the foreign keys currently defined on the table. To add a foreign key, click the Add button; to delete a foreign key, select it and click the Remove button.

Note:

After you add a foreign key, **to add another foreign key**, click the **Add** button.

Name: Name of the foreign key definition. Must be unique within the database.

Enabled: If this option is checked, the foreign key is enforced.

Referenced Schema: Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers.

Referenced Table: Name of the table with the primary key or unique constraint to which this foreign key refers.

Referenced Constraint: Name of the primary key or unique constraint to which this foreign key refers.

Associations: Local Column: Lists the column in the currently selected (local) table that is included in the foreign key definition. For each local column in the foreign key definition, select the name of a column in the local table.

Associations: Referenced Column on [table]: For each local column, identifies the column in the other (foreign) table that must have a value matching the value in the local column.

6.70.6 Check Constraints

Specifies one or more check constraints for the table. A check constraint specifies a condition that must be met when a row is inserted into the table or when an existing row is modified.

Check Constraints: Lists the check constraints currently defined on the table. To add a check constraint, click the Add button; to delete a check constraint, select it and click the Remove button.

Note:

After you add a check constraint, **to add another check constraint**, click the **Add** button.

Name: Name of the check constraint definition. Must be unique within the database.

Enabled: If this option is checked, the check constraint is enforced.

Condition: Condition that must be met for a row. Can be any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify: `rating >=1` and `rating <= 10`

To add another check constraint, click the **Add** button.

6.70.7 Indexes pane

Specifies properties for each index on the table.

Indexes: Lists the indexes currently defined on the table. (Note, however, that this display does *not* include any indexes related to primary key and unique key constraints.) To add an index, click the Add Index (+) icon; to delete an index, select it and click the Remove Index (X) icon.

Note:

After you add an index, **to add another index**, click the Add Index (+) icon.

Name: Name of the index. Must be unique within the schema.

Index: A list of index expressions, that is, the table columns or column expressions in the index. To add an index expression, click the Add Column Expression (+) icon; this adds a column name here and in Column Expression, where you can edit it. To delete an index expression, click the Remove Column Expression (X) icon; to move an index expression up or down in the list, click the Move Column Up and Move Column Down icons. An index must have at least one index expression.

For example, to create an index on the AUTHOR_LAST_NAME column of the BOOKS table from the tutorial (see [Create a Table \(BOOKS\)](#) (page 5-2)), click the + icon, and select AUTHOR_LAST_NAME in Column Name or Expression (next field), which changes BOOKS to AUTHOR_LAST_NAME in the Index field.

Column Name or Expression: A column name or column expression. A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.

Order: ASC for an ascending index (index values sorted in ascending order); DESC for a descending index (index values sorted in descending order).

6.70.8 In-Memory pane

Specifies In-Memory Column Store (IM column store) options for the table. The IM column store is an optional, static SGA pool that stores copies of tables and partitions in a special columnar format optimized for rapid scans. The IM column store does not replace the buffer cache, but acts as a supplement so that both memory areas can store

the same data in different formats. The IM column store is available starting with Oracle Database 12c Release 1 (12.1.0.2).

Note:

These IM column store features are *not* related to Oracle TimesTen In-Memory Database.

Define In-Memory Properties: Lets you define IM column store properties.

In-Memory: Enables or disables the IM column store feature for the table.

Compression (for table and for column clauses): The compression method for table data stored in the IM column store. (This data is called in-memory data.) Options: None (no compression), DML (optimized for DML operations and performs little or no data compression), QUERY LOW (least compression except for DML, resulting in the best query performance, QUERY HIGH (more compression than QUERY LOW but less than CAPACITY LOW), CAPACITY and CAPACITY LOW (more compression than QUERY HIGH, resulting in excellent query performance).

Priority: The data population priority for table data in the IM column store. Table data is populated in the IM column store before data for database objects with lower priority. The highest priority is `Critical`, and the lowest possible priority is `None`.

Distribute: Applicable only if you are using Oracle Real Application Clusters (Oracle RAC). It controls how table data in the IM column store is distributed across Oracle RAC instances.

Duplicate: Applicable only if you are using Oracle Real Application Clusters (Oracle RAC) on an engineered system. It controls how table data in the IM column store is duplicated across Oracle RAC instances.

Column Clauses: You can enable and disable specific table columns for the IM column store, and identify the data compression method for specific columns.

6.70.9 Storage pane

Lets you specify [Storage Options](#) (page 6-54) for the table.

Related Topics:

[Hive Tables: Storage](#) (page 6-54)

6.70.10 Column Sequences pane

Enables you to specify sequences and before-insert triggers to be used in populating a column with values. This approach is especially convenient for automatically populating primary key column values with unique values.

Column: Name of the column for which a sequence and a trigger are to be used to insert unique values. The data type of the column must be numeric.

Sequence: None causes no sequence to be used; Existing Sequence uses the sequence that you specify; New Sequence creates a new sequence with a default or specified name.

Trigger: Before-insert trigger that automatically inserts the next value of the specified sequence into the column when a new row is inserted.

6.70.11 Table Properties pane

Enables you to specify table properties such as compression, parallelism, and storage options.

Compress (heap-organized tables only): If this option is checked, data segments are compressed to reduce disk use. This clause is especially useful in environments such as data warehouses, where the amount of insert and update operations is small, and in OLTP environments.

Parallel: If this option is checked, parallel creation of the table is enabled, and the default degree of parallelism is set for queries and the DML INSERT, UPDATE, DELETE, and MERGE statements after table creation. You can also enter an integer in the text box to specify the degree of parallelism, which is the number of parallel threads used in the parallel operation. (Each parallel thread may use one or two parallel execution servers.) If you specify Parallel without entering an integer, the optimum degree of parallelism is automatically calculated.

Related Topics:

[Hive Tables: Table Properties](#) (page 6-54)

6.70.12 LOB Parameters pane

Specifies storage options for LOB (large object) columns, enabling you to override the default storage options.

Column: Name of the LOB column.

LOB Parameters: If this option is checked, the specified values for the remaining field are used. If this option is not checked, the default values for all fields are used.

Segment: LOB segment ID.

Tablespace: Name of the tablespace for the LOB data.

Store in Row: If this option is checked, the LOB value is stored in the row (inline) if its length is less than approximately 4000 bytes minus system control information.

Cache: Specifies how Oracle Database should store blocks in the buffer cache:

- **CACHE**: For data that is accessed frequently, indicates that the blocks retrieved for this table are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables.
- **NOCACHE**: For data that is not accessed frequently, indicates that the blocks retrieved for this table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed. NOCACHE is the default for LOB storage.
- **CACHE READS**: LOB values are brought into the buffer cache only during read operations but not during write operations.

Retention: If this option is checked, old versions of this LOB column are retained. You can specify this option only if the database is running in automatic undo mode and if you do not specify a Pct Version value.

Logging: <DEFAULT> means to use the Oracle Database default. ON means that the table creation and any subsequent direct loader (SQL*Loader) and direct-path INSERT

operations against the table, partition, or LOB storage are logged in the redo log file. OFF means that these operations are not logged in the redo log file.

Chunk: The number of bytes to be allocated for LOB manipulation. If the value is not a multiple of the database block size, then the database rounds up in bytes to the next multiple. The maximum value is 32768 (32K), which is the largest Oracle Database block size allowed. The default CHUNK size is one Oracle Database block.

Pct Version: Specifies the maximum percentage of overall LOB storage space used for maintaining old versions of the LOB. The default value is 10, meaning that older versions of the LOB data are not overwritten until they consume 10% of the overall LOB storage space. You can specify a Pct Version value whether the database is running in manual mode (where it is the default) or automatic undo mode (where Retention is the default). You cannot specify both a Pct Version value and the Retention option.

Free Pools: Specifies the number of groups of free lists for the LOB segment, usually the number of instances in a Real Application Clusters environment or 1 for a single-instance database. You can specify this option only if the database is running in automatic undo mode. You cannot specify both a Free Pools value and the Free Lists fields.

Extents - Initial: Size of the first extent of the table. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Extents - Next: Size of the next extent to be allocated to the table. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Extents - Min: Minimum number of extents allocated when the table is created.

Extents - Max: Maximum number of extents allocated when the table is created.

Unlimited (if checked) means that there is no maximum (and any specified maximum is ignored).

Extents - Pct Increase: Percentage that each extent grows over the previous extent.

Buffer Pool: <DEFAULT> means to use the Oracle Database default. KEEP means to put blocks from the segment into the Keep buffer pool; maintaining an appropriately sized Keep buffer pool lets Oracle retain the database object in memory to avoid I/O operations. RECYCLE means to put blocks from the segment into the Recycle pool; an appropriately sized Recycle pool reduces the number of objects whose default pool is the Recycle pool from taking up unnecessary cache space.

Free Lists: Number of free lists for each of the free list groups for the table. The default and minimum value for this parameter is 1, meaning that each free list group contains one free list.

Free List Groups: Number of groups of free lists for the table. The default and minimum value for this parameter is 1. Oracle uses the instance number of Real Application Clusters instances to map each instance to one free list group.

6.70.13 Partitioning pane

Specifies partitioning options for a **partitioned table**, which is a table that is organized into smaller and more manageable pieces called partitions. SQL queries and DML statements do not need to be modified in order to access partitioned tables; however, after partitions are defined, DDL statements can access and manipulate individual partitions rather than entire tables or indexes. Also, partitioning is entirely transparent to applications.

Partition By: The type of partitioning: **RANGE** partitions the table on ranges of values from the column list (which for an index-organized table must be a subset of the primary key columns of the table); **HASH** partitions the table using the hash method (rows assigned to partitions using a hash function on values found in columns designated as the partitioning key); **LIST** partitions the table on lists of literal values from column (useful for controlling how individual rows map to specific partitions).

Available: Lists the columns whose values are available to be used in assigning rows to partitions.

Selected: Lists the column whose values are to be used in assigning rows to partitions.

To add a column to the partitioning definition, select it in Available Columns and click the Add (>) icon; to remove a column from the partitioning definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the partitioning definition, select it in Selected Columns and use the arrow buttons.

Subpartition By: The partitioning type to be used to create subpartitions within each range partition. Use the Available and Selected column lists select and deselect a column for subpartitioning.

Related Topics:

[Hive Tables: Partitions](#) (page 6-54)

6.70.14 Subpartition Templates pane

Specifies subpartitioning options for a partitioned table. The options depend on the subpartition type, and might include the following.

Hash Quantity: Hash subpartition quantity.

Tablespaces: Available and Selected tablespaces for storage of the data in a subpartition.

Subpartition Templates: Specifications (subpartition templates) to control the placement of rows in each subpartition. Click the Add (+) icon to add a subpartition template that is appropriate for the subpartition type.

Subpartition Details: For each subpartition template, specify a name and (if relevant) a value or set of values that is appropriate for the subpartition type.

Storage: Enables you to specify a tablespace for the subpartition.

6.70.15 Partition Definitions pane

Defines each partition for a partitioned table. The options depend on the partition type, and might include the following.

Partitions: Specifications to control the placement of rows in each partition. Click the Add (+) icon to add a partition specification that is appropriate for the partition type.

Partition Details: For each partition specification, specify a name and (if relevant) a value or set of values that is appropriate for the subpartition type.

Storage: Enables you to specify a tablespace for the partition.

Subpartitions: Enables you to specify subpartition information.

6.70.16 Index Organized Properties pane

Specifies options for an **index-organized table**, which is a table in which the rows, both primary key column values and nonkey column values, are maintained in an index built on the primary key. Index-organized tables are best suited for primary key-based access and manipulation.

PCTTHRESHOLD: The percentage of space reserved in the index block for an index-organized table row; must be large enough to hold the primary key. All trailing columns of a row, starting with the column that causes the specified threshold to be exceeded, are stored in the overflow segment. PCTTHRESHOLD must be a value from 1 to 50; the default is 50.

Key Compression: If this option is checked, key compression is enabled, which eliminates repeated occurrence of primary key column values in index-organized tables. In the box to the right of this field, you can specify the prefix length, which is the number of prefix columns to compress. (This value can be from 1 to the number of primary key columns minus 1; the default prefix length is the number of primary key columns minus 1.)

Include Column: Column at which to divide an index-organized table row into index and overflow portions. The primary key columns are always stored in the index. The Include Column can be either the last primary key column or any non-primary-key column. All non-primary-key columns that follow the Include Column are stored in the overflow data segment.

Mapping Table: If this option is checked, SQL Developer creates a mapping of local to physical ROWIDs and store them in a heap-organized table. This mapping is needed in order to create a bitmap index on the index-organized table. If the index-organized table is partitioned, then the mapping table is also partitioned and its partitions have the same name and physical attributes as the base table partitions.

Overflow: Specifications for the overflow segment. The options are the same as in the [Storage Options](#) (page 6-54) dialog box.

6.70.17 External Table Properties pane

Specifies options for an **external table**, which is a read-only table whose metadata is stored in the database but whose data is stored outside the database. Among other capabilities, external tables enable you to query data without first loading it into the database.

Access Driver: The access driver of the external table. The access driver is the API that interprets the external data for the database: ORACLE_LOADER or ORACLE_DATAPUMP. You must specify the ORACLE_DATAPUMP access driver if you specify the AS subquery clause to unload data from one Oracle database and reload it into the same database or a different Oracle database.

Access Type: Type of data to be automatically converted during loads and unloads: BLOB or CLOB.

Default Directory: A default directory object corresponding to a directory on the file system where the external data sources may reside. The default directory can also be used by the access driver to store auxiliary files such as error logs.

Project Column: Determines how the access driver validates the rows of an external table in subsequent queries. **ALL** processes all column values, regardless of which columns are selected, and validates only those rows with fully valid column entries. If any column value would raise an error, such as a data type conversion error, the row

is rejected even if that column was not referenced in the select list. **REFERENCED** processes only those columns in the select list.

The ALL setting guarantees consistent result sets. The REFERENCED setting can result in different numbers of rows returned, depending on the columns referenced in subsequent queries, but is faster than the ALL setting. If a subsequent query selects all columns of the external table, then the settings behave identically.

Reject Limit: The number of conversion errors can occur during a query of the external data before an Oracle Database error is returned and the query is aborted.

Access Parameters: Values to the parameters of the specific access driver for this external table.

Location Specifications: One or more external data sources. Each is usually a file, but it need not be. Oracle Database does not interpret this clause; it is up to the access driver to interpret this information in the context of the external data. Use the Add (+) icon to add each location specification.

6.70.18 Comment pane

Optional descriptive comment about the table.

6.70.19 DDL pane

You can review and save the CREATE TABLE statement that SQL Developer will use to create a new table or that reflects any changes you have made to the table properties. If you want to make any changes, go back to the relevant tabs and make the changes there.

To save the SQL statement to a script file, click **Save** and specify the location and file name.

When you are finished, click **OK**.

Related Topics

- [Create Table \(quick creation\)](#) (page 6-41) (dialog box including advanced features)
- [Tables](#) (page 1-23)
- [Hive Tables](#) (page 6-53)

6.70.20 Hive Tables

For tables created using an Oracle Hive connection, the information is different for the following:

- [Hive Tables: Partitions](#) (page 6-54)
- [Hive Tables: Buckets](#) (page 6-54)
- [Hive Tables: Skew](#) (page 6-54)
- [Hive Tables: Storage](#) (page 6-54)
- [Hive Tables: Table Properties](#) (page 6-54)

6.70.20.1 Hive Tables: Partitions

Select one or more columns to use for partitioning, and under Partitions add one item for each partition. Use the tabs below Partitions to define the following information for each:

- **Details:** Specification details.
- **Location:** Partition location.
- **Storage Format:** Predefined file format, Input format class, and Output format class.
- **SerDe:** Java class for the SerDe (Serializer and Deserializer) interface, used to transform a string or binary record into a Java object that Hive can manipulate. For more information, see <http://blog.cloudera.com/blog/2012/12/how-to-use-a-serde-in-apache-hive/>.

6.70.20.2 Hive Tables: Buckets

Select one or more columns to use for bucket sorting, and under Bucket Sort Columns add one item for each combination of Expression and Order.

6.70.20.3 Hive Tables: Skew

Select one or more columns to use for creating a skewed table, and under Skewed Values add one item for each column to be used. Skewed tables can be used to improve performance for tables where one or more columns have skewed values (values that appear very often). For more information, see <https://cwiki.apache.org/confluence/display/Hive/LanguageManual+DDL#LanguageManualDDL-SkewedTables>.

6.70.20.4 Hive Tables: Storage

Define properties as needed for Native (Delimited or SerDe) or Non-Native storage types.

6.70.20.5 Hive Tables: Table Properties

Add any table properties (Name and Value for each) as needed.

Related Topics

- [Create Table \(quick creation\)](#) (page 6-41) (dialog box including advanced features)
- [Tables](#) (page 1-23)
- [Hive Tables](#) (page 6-53)

6.71 Storage Options

When you create or edit a table or an index, you can override the default storage options.

Tablespace: Name of the tablespace for the table or index.

Pct Free: Percentage of space in each of the data blocks of the table or index reserved for future updates. You can enter a value from 0 through 99.

Pct Used: Minimum percentage of used space that Oracle maintains for each data block. A block becomes a candidate for row insertions when its used space falls below the Pct Used value. You can enter a value from 1 through 99.

Logging: <DEFAULT> means to use the Oracle Database default. ON means that the table creation and any subsequent direct loader (SQL*Loader) and direct-path INSERT operations against the table, partition, or LOB storage are logged in the redo log file. OFF means that these operations are not logged in the redo log file.

Ini Trans: Number of update transaction entries for which space is initially reserved in the data block header.

Max Trans: Number of transaction entries that could concurrently use data in a data block. This parameter has been deprecated. Oracle Database now automatically allows up to 255 concurrent update transactions for any data block, depending on the available space in the block.

Extents - Initial: Size of the first extent of the table or index. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Extents - Next: Size of the next extent to be allocated to the table or index. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Extents - Min: Minimum number of extents allocated when the table or index is created.

Extents - Max: Maximum number of extents allocated when the table or index is created. **Unlimited** (if checked) means that there is no maximum (and any specified maximum is ignored).

Pct Increase: Percentage that each extent grows over the previous extent.

Buffer Pool: <DEFAULT> means to use the Oracle Database default. KEEP means to put blocks from the segment into the Keep buffer pool; maintaining an appropriately sized Keep buffer pool lets Oracle retain the database object in memory to avoid I/O operations. RECYCLE means to put blocks from the segment into the Recycle pool; an appropriately sized Recycle pool reduces the number of objects whose default pool is the Recycle pool from taking up unnecessary cache space.

Free Lists: Number of free lists for each of the free list groups for the table or index. The default and minimum value for this parameter is 1, meaning that each free list group contains one free list.

Free List Groups: Number of groups of free lists for the table or index. The default and minimum value for this parameter is 1. Oracle uses the instance number of Real Application Clusters instances to map each instance to one free list group.

Related Topics

- [Create/Edit Table \(with advanced options\)](#) (page 6-42) (dialog box including advanced features)
- [Create/Edit Index](#) (page 6-28) (dialog box)

6.72 Create/Edit Tablespace

This dialog box is used for creating a new tablespace or editing an existing tablespace. (Tablespaces appear on the DBA navigator under Storage.) The tablespace properties are grouped under several tabs.

Name: Name of the tablespace.

Tablespace Type: **Permanent** (contains persistent schema objects that are stored in data files), **Temporary** (contains schema objects only for the duration of a session, and stored in temp files), or **Undo** (a type of permanent tablespace used by Oracle Database to manage undo data if you are running your database in automatic undo management mode).

6.72.1 File Specifications tab

Specifies data files for the tablespace.

File Specifications: Contains one entry for each data file in the tablespace. To add a file, click the Add (plus sign) icon; to remove a file from the tablespace (but not delete the file itself), click the Remove (X) icon.

File Specification Properties: Displays properties of the selected data file.

Name: Name of the data file.

Directory: Name of the directory or folder for the data file. If not specified, the default location for tablespace data files is used.

File Size: Maximum size for the file; you can specify K (kilobytes) or M (megabytes) for the unit of measure.

Reuse Existing File: If this option is enabled and if a file with the same name already exists, the existing file is used. If this option is disabled and if a file with the same name already exists, a warning message is displayed.

Auto Extend: If this option is enabled, automatic extension is performed for a new or existing data file or temp file.

Next Size: The size in bytes of the next increment of disk space to be allocated automatically when more extents are required. The default is the size of one data block.

Max Size: The maximum disk space allowed for automatic extension of the data file.

6.72.2 Properties tab

Specifies general properties for the database role.

File Type: **Small File** (a traditional Oracle tablespace, which can contain 1022 data files or temp files, each of which can contain up to approximately 4 million (2^{22}) blocks) or **Big File** (contains only one data file or temp file, which can contain up to approximately 4 billion (2^{32}) blocks).

Online Status: **Online** (makes the tablespace available immediately after creation to users who have been granted access to the tablespace) or **Offline** (make the tablespace unavailable immediately after creation).

Block Size: You can specify a nonstandard block size for the tablespace. The integer you specify in this clause must correspond with the setting of one `DB_nK_CACHE_SIZE` parameter setting.

Logging: **Logging** (causes the creation of a database object, as well as subsequent inserts into the object, to be logged in the redo log file), **No Logging** (causes these operations not to be logged in the redo log file), or **File System Logging** (valid only for logging of SecureFiles LOB segments; specify this if you want to log only metadata changes).

Force Logging: If this option is enabled, Oracle Database will log all changes to all objects in the tablespace except changes to temporary segments, overriding any

NOLOGGING setting for individual objects. The database must be open and in READ WRITE mode.

Segment Management: Auto (the database manages the free space of segments in the tablespace using a bitmap; also called automatic segment-space management), or **Manual** (the database to manage the free space of segments in the tablespace using free lists).

Extent Management Type: Local (uses bitmaps in the tablespaces themselves to manage extents) or Dictionary (uses the data dictionary to manage extents).

Auto Allocate: If this option is enabled, Oracle Database determines the optimal size of additional extents, with a minimum extent size of 64 KB.

Uniform Size (available only if Auto Allocate is disabled): You can specify an extent size or use the default size of 1 MB. All extents in the tablespace are of this size. Locally managed temporary tablespaces can only use this type of allocation.

6.72.3 Default Parameters tab

Table Compression: Type of compression for table data in the tablespace: **No Compress** (no compression performed), **Compress** (basic table compression: Oracle Database attempts to compress data during direct-path INSERT operations when it is productive to do so), **Compress for All Operations** (OLTP table compression: Oracle Database compresses data during all DML operations on the table), or **Compress for Direct Load** (same as Compress).

Define Extents: Enables you to specify default values for extent-related parameters (potentially overriding any standard Oracle Database default values).

Initial Extent: Size of the first extent of the tablespace. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Next Extent: Size of the next extent to be allocated to the tablespace. Specify K (kilobytes) or M (megabytes) for the unit associated with the number.

Min Extents: Minimum number of extents allocated when the tablespace is created.

Max Extents: Maximum number of extents allocated when the tablespace is created.

Unlimited (if checked) means that there is no maximum (and any specified maximum is ignored).

Pct Increase: Percentage that each extent grows over the previous extent.

6.72.4 DDL tab

Displays the SQL statements that SQL Developer will use to create a new tablespace or to edit an existing tablespace. This display is read-only; if you want to make any changes, go back to the relevant tabs and make the changes there.

6.73 Create Trigger

The following information applies to a trigger, which is which is a stored PL/SQL block associated with a table, a schema, or the database, or an anonymous PL/SQL block or a call to a procedure implemented in PL/SQL or Java. The trigger is automatically executed when the specified conditions occur.

Schema: Database schema in which to create the trigger.

Name: Name of the trigger. Must be unique within the database.

Add New Source in Lowercase: If this option is checked, new text is entered in lowercase regardless of the case in which you type it. This option affects only the appearance of the code, because PL/SQL is not case-sensitive in its execution.

Trigger tab

Trigger Type: The type of object on which to create the trigger: TABLE, VIEW, SCHEMA, or DATABASE. (The remaining items depend on the type of trigger.)

Table Owner or View Owner: For a trigger on a table or a view, the name of the owner of the table or the view.

Table Name or View Name : For a trigger on a table or a view, the name of the table or the view.

Before or After: For a trigger on a table, select Before to cause the database to fire the trigger before executing the triggering event, or select After to cause the database to fire the trigger after executing the triggering event.

Statement Level or Row Level: For a trigger on a table, Statement Level fires the trigger once before or after the triggering statement that meets the optional trigger constraint defined in the WHEN condition; Row Level fires the trigger once *for each row* that is affected by the triggering statement and that meets the optional trigger constraint defined in the WHEN condition.

Insert, Update, Delete: For a trigger on a table or a view, Insert fires the trigger whenever an INSERT statement adds a row to a table or adds an element to a nested table; Update fires fire the trigger whenever an UPDATE statement changes a value in one of the columns specified in Selected Columns (or in any column if no columns are specified); Delete fires the trigger whenever a DELETE statement removes a row from the table or removes an element from a nested table.

Referencing - Old: For a trigger on a table, the correlation names in the PL/SQL block and WHEN condition of a row trigger to refer specifically to old value of the current row.

Referencing - New: For a trigger on a table, the correlation names in the PL/SQL block and WHEN condition of a row trigger to refer specifically to new value of the current row.

Available Columns: For a trigger on a table, lists the columns from which you can select for use in an Update trigger definition.

Selected Columns: For a trigger on a table, lists the columns used in an Update trigger definition.

When: For a trigger on a table, an optional trigger condition, which is a SQL condition that must be satisfied for the database to fire the trigger. This condition must contain correlation names and cannot contain a query.

Schema: For a trigger on a schema, the name of the schema on which to create the trigger.

Available Events: For a trigger on a schema or database, lists events from which you can select for use in the trigger definition.

Selected Events: For a trigger on a schema or database, lists events used in the trigger definition.

DDL Tab

You can view a SQL CREATE statement that reflects the current definition of the object, or a SQL ALTER statement to modify an existing object to reflect your changes.

Related Topics

- [Triggers](#) (page 1-25)

6.74 Create Type (User-Defined)

This dialog box is displayed when you right-click Types in the Connections navigator and select Create Type to create a user-defined type. After you complete the information in this dialog box and click OK, a SQL Worksheet is displayed in which you must specify the appropriate definition of the type.

Schema: Database schema in which to create the type.

Name: Name of the type. Must be unique within its schema.

Type: Select the type of data type to be created: array type, object type specification, object type specification and type body, or table type.

For more information about creating a user-defined type, see the CREATE TYPE statement in *Oracle Database SQL Language Reference*.

Related Topics

- [Types](#) (page 1-26)

6.75 Create/Edit User

This dialog box is used for creating a new database user or editing an existing database user. The user properties are grouped under several tabs.

To create or edit a database user, the user associated with your database connection must have the DBA role. You should also be familiar with the main concepts and techniques documented in *Oracle Database Administrator's Guide* and in the chapter about managing security for Oracle Database users in *Oracle Database Security Guide*.

6.75.1 User tab

Specifies general properties for the database user.

User Name: The user name string. For an existing user, this field is read-only; to change the name, you must drop the user and create a new user with the desired name.

New Password: Password string for the new user, or new password for an existing user. You must also type the same password string for **Confirm Password**.

Password Expired: If this option is checked, the password is marked as expired, and the user must change the password before being permitted to connect to the database.

Account Locked: If this option is checked, the user will not be permitted to connect to the database until a DBA user unlocks the account associated with this user.

6.75.2 Granted Roles tab

Specifies roles to be granted to the user. For each role, you can check **Granted** to grant the role, **Admin** to permit the user to grant the role to other users, and **Default** to use the default settings for Granted and Admin.

For convenience, you can click buttons to affect all settings (Grant All, Revoke All, Admin All, Admin None, Default All, Default None); then, you can specify other settings for individual roles.

6.75.3 System Privileges tab

Specifies privileges to be granted to the user. For each privilege, you can check **Granted** to grant the privilege, and **Admin** to permit the user to grant the privilege to other users.

For convenience, you can click buttons to affect all settings (Grant All, Revoke All, Admin All, Admin None); then, you can specify other settings for individual privileges.

6.75.4 Quotas tab

Specifies disk usage limits on specified tablespaces for the user. If you check Unlimited, there is no disk usage limit on the tablespace.

6.75.5 SQL tab

Displays the SQL statements that SQL Developer will use to create (after executing a CREATE USER statement) a new user or to edit an existing user. This display is read-only; if you want to make any changes, go back to the relevant tabs and make the changes there.

Related Topics

- [Users \(Other Users\)](#) (page 1-26)

6.76 Create/Edit User Defined Report

The following information applies to a user-defined report. You can create a simple report based on a SQL query or, if you click Advanced, a report with specialized features such as a child report.

Advanced: If this option is not checked, the box contains basic information and the Property subtab. If this option is checked, the box also contains subtabs for SQL Query, Child Reports, Binds, Frill Down, and PDF.

Database connection dropdown: Select the database connection to use for the report.

Style: The style or type of report:

- **Table:** Tabular presentation (default)
- **Chart:** Bar or pie chart
- **Gauge:** Dial or status meter
- **Code:** Formatted code
- **Script:** Executable script
- **PL/SQL DBMS Output:** Output in PL/SQL DBMS format

You are encouraged to experiment with different style options to see which best meets your needs.

Description: Optional description of the report.

Tool Tip: Popup text that will be displayed when you hover over the report name in the Reports navigator.

Test: Tests the report definition by running it in a separate window. This feature enables you to test the report before creating it.

SQL: Area for entering the SQL query for the report. You can use the folder icon to select a file with the SQL text, and the eraser icon to erase the contents of the SQL box.

SQL Query (advanced)

Area for entering the SQL query for the report. You can use the folder icon to select a file with the SQL text, and the eraser icon to erase the contents of the SQL box.

SQL Variants: Lets you add and remove Minimum Version and Maximum Version information.

Child Reports (advanced)

Lets you add and remove child reports to a master report.

For an example of creating a child report (where you can click a department ID in the top part to display employees in that department in the bottom part), go to the **Start Page** (Help > Start Page). On the Get Started tab under Featured Online Demonstrations, and click **Reporting Features**. On the displayed page, find the title containing **Adding a Child Report**.

Add Child: Lets you specify the name of the child report, after which a node is added for that child report, along with SQL Query, Property, and PDF subtabs for the child report.

Binds (advanced)

Name: Name of the bind variable.

Prompt: String displayed when the user is prompted to enter a value. Example: *Table name*

Default: Default value if the user does not enter a value at the prompt. To accept the Oracle SQL value, specify NULL_VALUE.

ToolTip: Optional tooltip text to be displayed when the mouse pointer stays briefly over the bind variable name.

Drill Down (advanced)

You can click **Add Report** to specify one or more reports to be enabled for drill-down display. To remove a report from drill-down, select it and click **Remove Report**.

Property

Displays options relevant to the Style (report type) that you selected. (Relevant for Table, Chart, and Gauge styles.)

PDF (advanced)

Lets you specify layout and formatting options for when the report is exported as PDF. For example, you can specify a header and footer; security options including password-protection, encryption, and restrictions on printing and copying; and page layout options such as margin sizes and portrait or landscape orientation.

Related Topics

- [SQL Developer Reports](#) (page 1-82)
- [User Defined reports](#) (page 1-89)
- [Gauges: In the SQL Worksheet and User-Defined Reports](#) (page 1-49)

6.77 Create/Edit User Defined Report Folder

The following information applies to a folder for organizing user-defined reports. Each folder can contain reports and other folders (subfolders). For example, you can create a folder named Sales, and then under that folder create folders named Sales by District and Sales by Product.

Name: Name of the folder.

Description: Optional description of the folder.

ToolTip: Optional tooltip text to be displayed when the mouse pointer stays briefly over the folder name in the Reports navigator display.

Related Topics

- [SQL Developer Reports](#) (page 1-82)
- [User Defined reports](#) (page 1-89)

6.78 Create/Edit View

The view dialog box is used for creating or editing a view or materialized view. You can use the SQL Query tab or a series of panes to specify the query part of the view definition, and you can use one or more other panes (depending on the type of view) for other parts of the definition.

For Edit View, you cannot change the definition of an existing view, but you can define, modify, or drop view constraints.

If you click OK before you are finished creating or editing the view, right-click the view name in the Connections navigator, select Edit, and continue creating or editing the view.

Schema: Database schema in which to create the view.

Name: Name of the view. Must be unique within a schema.

Advanced: If this option is checked, the dialog box changes to include a pane that provides an extended set of features for creating the view.

6.78.1 SQL Query tab or pane

As a tab (if you did not check the Advanced box), it contains the SQL code for the query part of the view definition, using the SELECT and FROM keywords and usually a WHERE clause with whatever syntax is needed to retrieve the desired information.

As a pane (if you checked the Advanced box), it presents options for building specific parts of the query.

For example, the following query, from the [Create a View](#) (page 5-8) tutorial topic, selects columns from the PATRONS and TRANSACTIONS tables, ordering them first by values in the PATRON_ID column in the PATRONS table and then by values in the TRANSACTION_TYPE column in the TRANSACTIONS table. The result is a listing by patron ID of all patrons who had transactions, and for each listed patron the transaction information listed by transaction type

```
CREATE VIEW patrons_trans_view AS
  SELECT p.patron_id, p.last_name, p.first_name,
         t.transaction_type, t.transaction_date
  FROM patrons p, transactions t
```



```
WHERE p.patron_id = t.patron_id
ORDER BY p.patron_id, t.transaction_type;
```

SQL Parse Results: If you click Test Syntax, displays any SQL syntax errors, or displays a message indicating no errors if there are no syntax errors.

Revert: Cancels any edits you have made in the SQL Query box, and displays the contents of the box before these edits.

Test Syntax: Checks the statement in the SQL Query box for any SQL syntax errors.

Test Query: Displays a dialog box that runs the SQL query and indicates the result. If there is an error in the query, such as an invalid character or a missing expression, the error is displayed.

6.78.2 View Properties or Materialized View Properties pane

Options for a standard view:

Force on create: If this option is checked, the view is created even if it has errors in its definition. This option is useful if you want to create the view regardless of any errors, and go back and correct the errors later. If this option is not checked, the view is not created if its definition contains any errors.

Query Restriction: If this option is checked, you can enable one of the following options

- **Read Only:** Prevents the view from being used to add, delete, or change data in the underlying table or tables.
- **Check Option:** If this option is checked, it prohibits any changes to the underlying table or tables that would produce rows that are not included in this view.

Options for a materialized view:

Refresh Options:

Method: The method of refresh operation to be performed:

- **Complete Refresh:** Executes the defining query of the materialized view, even if a fast refresh is possible.
- **Fast Refresh:** Uses the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes for conventional DML changes are stored in the materialized view log associated with the master table. The changes for direct-path INSERT operations are stored in the direct loader log.
- **Force Refresh:** Performs a fast refresh if one is possible; otherwise, performs a complete refresh.
- **Never:** Do not perform refresh operations.

When: The type of refresh operation to be performed:

- **On Demand:** Performs a refresh when one of the DBMS_MVIEW refresh procedures is called.
- **On Commit:** Performs a fast refresh whenever the database commits a transaction that operates on a master table of the materialized view. This may increase the time taken to complete the commit, because the database performs the refresh operation as part of the commit process.

- **Specify:** Performs refresh operations according to what you specify in the *Start on* and *Next* fields.
- **Never:** Does not perform a refresh operation.

Type: Refresh type, which determines the type of materialized view:

- **Primary Key:** Creates a primary key materialized view, which allows materialized view master tables to be reorganized without affecting the eligibility of the materialized view for fast refresh.
- **Row ID:** Creates a rowid materialized view, which is useful if the materialized view does not include all primary key columns of the master tables.

Start on: Starting date and time for the first automatic refresh operation. Must be in the future.

Next: Time for the next automatic refresh operation. The interval between the *Start on* and *Next* times establishes the interval for subsequent automatic refresh operations. If you do not specify a value, the refresh operation is performed only once at the time specified for *Start on*.

Constraints: If this option is checked, more rewrite alternatives can be used during the refresh operation, resulting in more efficient refresh execution. The behavior of this option is affected by whether you select *Enforced* or *Trusted*.

Enforced: Causes only enforced constraints to be used during the refresh operation.

Trusted: Enables the use of dimension and constraint information that has been declared trustworthy by the database administrator but that has not been validated by the database. If the dimension and constraint information is valid, performance may improve. However, if this information is invalid, then the refresh procedure may corrupt the materialized view even though it returns a success status.

Materialized View Options:

Parallel: If this option is checked, parallel operations will be supported for the materialized view, and you can specify a number for the default degree of parallelism for queries and DML on the materialized view after creation.

Enable Cache: If this option is checked, the blocks retrieved for this table are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table scan is performed. This setting is useful for small lookup tables. If this option is not checked, the blocks are placed at the least recently used end of the LRU list.

Build Type: Specifies when to populate the materialized view. **Immediate** indicates that the materialized view is to be populated immediately. **Deferred** indicates that the materialized view is to be populated by the next refresh operation. If you specify *Deferred*, the first (deferred) refresh must always be a complete refresh; until then, the materialized view has a staleness value of unusable, so it cannot be used for query rewrite.

Enable Query Rewrite: If this option is checked, the materialized view is enabled for query rewrite, an optimization technique that transforms a user request written in terms of master tables into a semantically equivalent request that includes one or more materialized views.

Prebuilt Option: If this option is checked, an existing table is registered as a preinitialized materialized view. This option is particularly useful for registering large materialized views in a data warehousing environment. The table must have the same name and be in the same schema as the resulting materialized view, and the table

should reflect the materialization of a subquery. **Reduced Precision** authorizes the loss of precision that will result if the precision of the table or materialized view columns do not exactly match the precision returned by subquery. **No Reduced Precision** requires that the precision of the table or materialized view columns match exactly the precision returned by subquery, or the create operation will fail.

Index Storage Options:

Use Index: If this option is checked, a default index is created and used to speed up incremental (fast) refresh of the materialized view. If this option is not checked, this default index is not created. (For example, you might choose to suppress the index creation now and to create such an index explicitly later.)

Use Tablespace: If this option is checked, you can specify the tablespace in which the materialized view is to be created. If this option is not checked, the materialized view is created in the default tablespace of the schema containing the materialized view.

6.78.3 DDL tab or pane

You can view a SQL CREATE statement that reflects the current definition of the object, or a SQL ALTER statement to modify an existing object to reflect your changes.

To save the SQL statement to a script file, click **Save** and specify the location and file name.

Related Topics

- [Views](#) (page 1-26)

6.78.4 Quick-Pick Objects pane

Specifies objects that you can use in the SELECT, FROM, and WHERE clauses of the view definition. Identify the tables and views on which this view is based, and the columns in those tables and views that are used in the definition of this view. To see the results of your quick-pick specification, either check Auto-Query or click Query.

Schema: Database schema containing the objects to be selected.

Type Filter - Filter Types: Enables you to limit the display of objects available for selection to certain types of database objects (for example, to show only tables or views).

Name Filter: Enables you to limit the display of objects available for selection according to a character string in the name, with the percent sign (%) as a wildcard character. For example, to limit the display of available tables and views to those whose names start with the string EM, specify the following name filter: EM%

Auto-Query: If this option is enabled, the display of available objects is automatically refreshed when you specify or change the Type Filter or Name Filter value.

Query: Refreshes the display of available objects based on the Type Filter and Name Filter values.

Available: Lists the objects (typically, tables and views in a hierarchical display) from which you can select objects to use in the SELECT, FROM, and WHERE clauses of the view definition.

Selected: Lists the objects (typically, columns) that you can use in the SELECT, FROM, and WHERE clauses of the view definition.

To add an object as selected, select it in Available and click the Add (>) icon; to remove an object as selected, select it in Selected and click the Remove (<) icon. To move all

objects from selected to available, use the Remove All (<<) icon. To move an object up or down in the selected list, select it in Selected and use the arrow buttons.

For the example in [DDL tab or pane](#) (page 6-65), select the DEPTNO and SAL columns from the EMP table.

6.78.5 FROM Clause pane

Specifies the tables and views that you can use in the FROM clause of the view definition.

Type Filter - Filter Types: Indicates whether or not you have limited the types of database objects to be displayed in the Available List (by clicking the **Filter Types** and selecting any combination of Materialized Views, Tables, and Views).

Naming Filter: Substring for limiting object names to appear as available. For example, to display only objects with names that start with EM, specify EM% (with the percent sign as a wildcard character).

Auto-Query: If this option is enabled, the display of available objects is automatically refreshed when you specify or change the Type Filter or Name Filter value.

Query: Refreshes the display of available objects based on the Type Filter and Name Filter values.

Available: Lists the tables and views that are available to be selected for use in the FROM clause of the view definition.

Selected: Lists the tables and views that you can use in the FROM clause of the view definition.

To add an object as selected, select it in Available and click the Add (>) icon; to remove an object as selected, select it in Selected and click the Remove (<) icon. To move all objects from available to selected, use the Add All (<<) icon; to move all objects from selected to available, use the Remove All (<<) icon.

If you attempt to remove any objects that have dependencies in the SQL query, a Confirm Remove box warns you that the dependencies will be removed with the FROM expression; to cancel the remove operation, click **No**.

Alias: Alias for the table or view.

For the example in [DDL tab or pane](#) (page 6-65), select the EMP table.

Join: If you select two tables and click this button, the [Edit Join](#) (page 6-90) dialog box is displayed.

6.78.6 SELECT Clause pane

Specifies objects that you can use in the SELECT clause of the view definition.

SELECT List: Lists the objects (typically, columns) that you can currently use in the SELECT clause. To add an object, click the Add (+) icon; to delete an object, select it and click the Delete (X) icon; to move an object up or down in the view definition, select it and use the up-arrow and down-arrow buttons.

Note:

After you add an object, **to add another object**, click the Add (+) icon.

Expression: Column name or an expression. For expressions, you can type them, or you can use the **Expression Palette** to add object names and function names.

Validate: Checks the validity of the Expression entry.

For the example in [DDL tab or pane](#) (page 6-65), select DEPTNO column and the MIN(emp.sal) and MAX(emp.sal) functions.

6.78.7 WHERE Clause pane

Specifies the WHERE clause of the view definition.

WHERE: The text of the WHERE clause, without the WHERE keyword. You can type the text completely; or you can type some of the text and use the **Expression Palette** to add object names, function names, and operators.

Example (from the [Create a View](#) (page 5-8) tutorial exercise): p.patron_id = t.patron_id

6.78.8 GROUP BY Clause pane

Specifies a clause to be used to group the selected rows based on the value of columns for each row and return a single row of summary information for each group. The GROUP BY clause groups rows but does not guarantee the order of the result set; to order the groupings, use the ORDER BY clause.

Available: Lists the tables and views, and the columns in each, that are available to be selected for use in the GROUP BY clause of the view definition.

Selected: Lists the tables and views, and the columns in each, that you can use in the GROUP BY clause of the view definition.

To add an object as selected, select it in Available and click the Add (>) icon; to remove an object as selected, select it in Selected and click the Remove (<) icon. To move all objects from available to selected, use the Add All (<<) icon; to move all objects from selected to available, use the Remove All (>>) icon.

6.78.9 HAVING Clause pane

Specifies an expression that must be satisfied for rows to be processed by the GROUP BY clause. For example, HAVING MIN(salary) < 30000 causes the GROUP BY clause to consider only rows where the minimum value of the relevant salary values is less than 30000.

HAVING: You can type the complete expression text, or you can use the **Expression Palette** to add object names, function names, and operators to the expression text.

6.78.10 ORDER BY Clause pane

Specifies one or more columns or column expressions whose values will be used to sort the results returned by the view. Without an ORDER BY clause, no guarantee exists that the same query executed more than once will retrieve rows in the same order.

ORDER BY List: Lists the objects (typically, columns) that you can currently use in the ORDER BY clause. To add an object, click the Add (+) icon; to delete an object, select it and click the Delete (X) icon; to move an object up or down in the view definition, select it and use the up-arrow and down-arrow buttons.

Note:

After you add an object, **to add another object**, click the Add (+) icon.

ORDER BY Expression Filter: For each column or column expression, you can type the text completely into the **Expression** box; or you can type some of the text and use the **Expression Palette** to add object names, function names, and operators.

Validate: Tests the validity of the syntax for the expression.

Order: ASC for ascending (expression values sorted in ascending order); DESC for descending (expression values sorted in descending order).

Nulls Ordering: NULLS FIRST to have null expression values appear before non-null values; NULLS LAST to have null expression values appear after non-null values. ("Before" and "after" positions are determined by the Order value.)

6.79 Create Window

This interface is used for creating a new Oracle Scheduler window.

(To create a window, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_WINDOW procedure.)

Name: Name of the window. It can optionally be prefixed with a schema name. Cannot be null. It is converted to uppercase unless enclosed in double-quotes.

Enabled: Enables the window. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the window is created.)

Resource Plan: The resource plan that automatically activates when the window opens. When the window closes, the system switches to the appropriate resource plan, which is usually the plan that was in effect before the window opened, but can also be the plan of a different window. Only one resource plan can be associated with a window. If null, the resource plan in effect when the window opens stays in effect for the duration of the window. If an empty string, the resource manager is disabled for the duration of the window. If the window is open and the resource plan is dropped, then the resource allocation for the duration of the window is not affected.

Duration: The length of time that the window stays open. Can range from one minute to 99 days.

Priority: Relevant when two windows overlap. Because only one window can be in effect at one time, the window priority determines which window opens. The two possible values for this attribute are HIGH and LOW. A high priority window has precedence over a low priority window, therefore, the low priority window does not open if it overlaps a high priority window.

Description: Optional descriptive text.

When to Open: Repeating or Schedule:

- For Repeating, specify the **Repeat Interval** and optionally the **Start Date** and **End Date**.
- For Schedule, select the name of the **Schedule** to be used.

Related Topics

- [Windows](#) (page 1-27)

- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.80 Create Window Group

This interface is used for creating a new Oracle Scheduler window group.

(To create a window group, SQL Developer internally uses the DBMS_SCHEDULER.CREATE_GROUP procedure.)

Name: Name of the window group. It can optionally be prefixed with a schema name. Cannot be null. It is converted to uppercase unless enclosed in double-quotes.

Enabled: Enables the window group. (Causes the DBMS_SCHEDULER.ENABLE procedure to be called after the window group is created.)

Comments: Optional comment text.

Members

Available Members: Lists all Scheduler windows.

Selected Members: List windows to be added to the window group. Use the arrow icons to move selected windows or all windows from one list to the other.

Related Topics

- [Window Groups](#) (page 1-27)
- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.81 Create XML Document from XML Schema

Use to specify the filename and the directory in which the XML file from an XML Schema is created, and options for the creation.

File Location

XML File: Name of the XML file to be created.

Directory: Directory for the file. You can click Browse to select a directory.

Use Registered Schemas: Select to use the schemas registered with SQL Developer. To view all the currently registered XML schemas, add new schemas to support additional namespaces and elements, remove existing schemas, and unload schemas from memory, select Tools > Preferences, XML Schemas.

Use File System Schema: Enter or browse to select the directory for the XML schema.

Options

Use to specify the default XML namespace and other options in creating the instance document.

Target Namespace: Select the XML namespace (xmlns attribute). Available namespaces are either schemas already registered with SQL Developer or specified in the file system from the previous wizard page.

Root Element: Select the root element. SQL Developer will automatically provide the root element from the schema selected as the target namespace.

Depth: Number of levels to display in the XML tree structure.

Encoding: Select the encoding that the XML parser will use to understand foreign characters.

Generate Only Required Elements: Select to generate XML elements with a required attribute in the instance document.

6.82 Create XML Schema

This dialog box enables you to create an XML schema object that can be associated with XML document instances.

Name: Name of the XML schema object.

After you click **OK**, an editing pane is displayed for specifying the URL of the XML schema.

6.83 Custom Filters

This dialog box is displayed if you right-click and select Customize Filters in the History tab for a .sql file or a SQL Trace (.trc) file. You can modify an existing filter or create a new filter.

Filter List: Names of the available filters. To edit an existing filter, select its name; the details for that filter are displayed in the dialog box. To remove an existing filter, select its name and click **Remove**.

To create a new filter, click **Add** and specify the filter name.

To restore the filters to those at SQL Developer installation, click **Restore Defaults**. This deleted any filters that have been added since installation.

Simple Expression: Create the filter by selecting a column and operator and by specifying a value. To select a value from a list, click **Insert**.

Complex Expression: Create the filter by entering a complex expression.

6.84 Customize Toolbars

This dialog box is displayed if you click View, Show Toolbars, Customize Main Toolbar, then Customize.

To add a command icon to the main toolbar, drag its icon from the dialog box to the desired location on the main toolbar.

To remove a command icon from the main toolbar, drag its icon from the main toolbar to the dialog box.

New Toolbar: Lets you add a "placeholder" in the main toolbar, into which you can drag icons.

Reset Toolbars: Undoes any main toolbar customizations, restoring it to the content when SQL Developer was originally installed. (To reset the entire user interface to its original form, see [Restoring the Original "Look and Feel"](#) (page 1-12).)

6.85 Data Pump Export

This wizard is displayed if you expand a connection in the DBA navigator, right-click Export Jobs, and select Data Pump Export Wizard.

To use the Data Pump Export wizard, you must understand the concepts and techniques for Oracle Data Pump, as documented in *Oracle Database Utilities*.

Related Topics:

[Using DBA Features in SQL Developer](#) (page 1-54)

6.85.1 Source page

Connection: Database connection to use to export the specified information for the specified types of objects.

Data or DDL: Specify whether to export only the data (applies to tables and views), only the object definitions (the "DDL"), or both the data and the object definitions.

Types to be Exported: Specify whether to export the database or just one or more tablespaces, schemas, or tables. (Your selection will affect the display of some remaining wizard pages.)

6.85.2 Tablespaces page

If you specified Tablespaces for Types to be Exported, select one or more tablespaces to be exported, and use the arrow icons to move them to the Selected Source Tablespaces box.

6.85.3 Schemas page

If you specified Schemas for Types to be Exported, select one or more schemas to be exported, and use the arrow icons to move them to the Selected Source Schemas box.

6.85.4 Tables page

If you specified Tables for Types to be Exported, select one or more tables to be exported, and use the arrow icons to move them to the Selected Source Schemas box.

Tables tab: To select one or more tables, see [Specify Objects](#) (page 6-137) under Shared Wizard pages.

Include/Exclude Filter tab: Optionally, filter the objects to be included or excluded, using the same interface as for the [Filter page](#) (page 6-71).

6.85.5 Filter page

Lets you specify data and metadata filtering to limit the type of information that is exported. You can specify multiple filters by adding rows.

If you specify any **Include** or **Exclude** rows, you must put any object names in single quotes. Example: **Filter Type** = TABLE, and **Expression** = not in ('EMPLOYEES')

6.85.6 Table Data page

See [Data](#) (page 6-137) under Shared Wizard Pages.

6.85.7 Options page

Lets you specify Data Pump Export options for the job.

6.85.8 Output Files page

Lets you specify one of more output (dump) files and certain options, including whether to delete existing dump files before creating new ones and whether to include the timestamp in dump file names.

6.85.9 Job Schedule page

Lets you specify scheduling options for the Data Pump Export job. The options include whether to start the job immediately or later, and any repeat options.

6.85.10 Summary page

See [Summary](#) (page 6-138) under Shared Wizard Pages.

On the Summary page, you can click the PL/SQL tab to see the PL/SQL procedure that will perform the job.

6.86 Data Pump Import

This wizard is displayed if you expand a connection in the DBA navigator, right-click Import Jobs, and select Data Pump Import Wizard.

To use the Data Pump Import wizard, you must understand the concepts and techniques for Oracle Data Pump.

Related Topics:

[Using DBA Features in SQL Developer](#) (page 1-54)

6.86.1 Type page

Connection: Database connection to use to import the specified information for the specified types of objects.

Job Name: Name to be given to this import job.

Data or DDL: Specify whether to import only the data (applies to tables and views), only the object definitions (the "DDL"), or both the data and the object definitions.

Types to be Imported: Specify whether to import the database from the dump file or files, or just one or more tablespaces, schemas, or tables. (Your selection will affect the display of some remaining wizard pages.)

Choose Import Files: Specify one or more dump files (that were created using Data Pump Export) from which to import the objects. To add files, click **Add Row**.

6.86.2 Tablespaces page

If you specified Tablespaces for Types to be Imported, select one or more tablespaces to be imported, and use the arrow icons to move them to the Selected Source Tablespaces box.

6.86.3 Schemas page

If you specified Schemas for Types to be Imported, select one or more schemas to be imported, and use the arrow icons to move them to the Selected Source Schemas box.

6.86.4 Tables page

If you specified Tables for Types to be Imported, select one or more tables to be imported, and use the arrow icons to move them to the Selected Source Schemas box.

Tables tab: To select one or more tables, see [Specify Objects](#) (page 6-137) under [Shared Wizard Pages](#) (page 6-137).

Include/Exclude Filter tab: Optionally, filter the objects to be included or excluded, using the same interface as for the [Filter page](#) (page 6-73).

6.86.5 Filter page

Lets you specify data and metadata filtering to limit the type of information that is imported. You can specify multiple filters by adding rows.

6.86.6 Table Data page

See [Data](#) (page 6-137) under [Shared Wizard Pages](#) (page 6-137).

6.86.7 Remapping page

Lets you specify options for the relevant Data Pump Import "remapping" parameters (REMAP_DATAFILE, REMAP_SCHEMA, REMAP_TABLESPACE).

6.86.8 Options page

Lets you specify Data Pump Import options for the job. The specific options depend on selections made on preceding wizard pages. Various options are implemented using DBMS_DATAPUMP subprograms, documented in *Oracle Database PL/SQL Packages and Types Reference*.

Note that the **Regenerate Object Ids** option calls the DBMS_DATAPUMP.the METADATA_TRANSFORM procedure and specifies OID as 0 (zero), which inhibits the assignment of the exported OID during type or table creation; instead, a new OID will be assigned.

6.86.9 Job Schedule page

Lets you specify scheduling options for the Data Pump Import job. The options include whether to start the job immediately or later, and any repeat options.

6.86.10 Summary page

See [Summary](#) (page 6-138) under [Shared Wizard Pages](#) (page 6-137).

On the Summary page, you can click the PL/SQL tab to see the PL/SQL procedure that will perform the job. For information about the Data Pump PL/SQL API, see *Oracle Database Utilities*.

6.87 Database Copy

This wizard is displayed if you click Tools, then Database Copy. You can copy objects, schemas, or a tablespace from a source connection to a destination connection. The specific options depend on what you are copying.

Some pages include **Proceed to Summary**, which lets you accept the default values for remaining pages and go directly to the Summary page by clicking **Next**.

6.87.1 Source/Destination page

Source Connection: Database connection from which to copy the objects, schemas, or tablespace.

Destination Connection: Database connection into which to copy the objects, schemas, or tablespace.

Note: For Schemas Copy or Tablespace Copy, the Source Connection user and the Destination Connection user must have DBA privileges. (DBA privileges are not required for Objects Copy.)

Copy Options: Specify whether to copy database objects, schemas, or tablespaces.

Objects Copy: Copies types of database objects and specific objects of those types.

Schemas Copy: Copies one or more schemas. If you select multiple schemas from the source connection, you can maintain the individual schemas in the destination connection or consolidate them all into one schema in the destination connection.

Tablespace Copy: Copies the objects from a source tablespace to a destination tablespace.

Copy DDL: Copies the object definitions. For destination objects with the same names as source objects of the same type (for example, if both contain a table named EMPLOYEES), specify whether to not perform the copy (that is, do *not* replace the destination objects) or to perform the copy (that is, replace the existing destination objects with the source objects).

Copy Data: Copies the data for any tables and views that are copied. If you do not select this option, any copied tables or views are empty in the destination connection.

Truncate Destination Data Before Copying: If a table or view of the same name already exists in the destination connection (for example, if an EMPLOYEES table exists in the source and destination), this option deletes any existing data in the destination object before copying the data from the source connection. If this option is not selected, the copied data is appended to the existing data in the destination object.

6.87.2 Object Types page

See [Types](#) (page 6-137) under Shared Wizard Pages.

6.87.3 Select Schemas page

If you specified Schemas Copy, select one or more schemas from the **Available Source Schemas** list and click the right-arrow to move them to **Selected Source Schemas**.

6.87.4 Select Tablespace page

If you specified Tablespace Copy, select the **Source Tablespace** and **Destination Tablespace**. Objects in the source tablespace are copied into the destination tablespace.

6.87.5 Specify Objects page

See [Specify Objects](#) (page 6-137) under Shared Wizard Pages.

6.87.6 Specify Data page

See [Data](#) (page 6-137) under Shared Wizard Pages.

6.87.7 Summary page

See [Summary](#) (page 6-138) under Shared Wizard Pages.

Related Topics

- [Database Connections](#) (page 1-28)

- [Database Objects](#) (page 1-13)

6.88 Database Differences

This wizard is displayed if you click Tools, then Database Diff. You can find differences between objects of the same type and name (for example, tables named CUSTOMERS) accessible from two different connections, and optionally generate DDL to update the objects in the destination schema to reflect the differences in the source schema.

Use the [Source/Destination page](#) (page 6-75) to specify the source and destination database connections. Database objects in the schemas associated with these connections will be compared. The schemas for the source and destination connections can be in the same database or different databases.

Some pages include **Proceed to Summary**, which lets you accept the default values for remaining pages and go directly to the Summary page by clicking **Next**.

6.88.1 Source/Destination page

Source Connection: Database connection that has access to the source objects to be compared.

Destination Connection: Database connection that has access to the destination objects to be compared.

DDL Generation Options: Check the options or features to be included if you later generate a script to update the objects in the destination schema to reflect differences in the source schema: comments, schema name (prefixed to schema objects), constraints, password values, constraint indexes, referential constraints. Options that are not checked are ignored or not performed during the script generation.

Schema: This option, if you also select Consolidate Schema under DDL Comparison Options, causes object names to be prefixed with the *target* schema name in the generated script (for example, CREATE TABLE "HR"."SKILLS" instead of CREATE TABLE "SKILLS"). If you select Maintain Schema under DDL Comparison Options, this Schema option is ignored and object names are prefixed with the *source* schema name in the generated script.

DDL Comparison Options: Check various options, including differences to be ignored, when objects in the source and destination schemas are compared.

Note:

The DDL Generation Options apply to how the DDL will be generated, whereas the DDL Comparison Options apply to how the comparison is made between the two objects.

Consolidate Schema and **Maintain Schema** enable you either to compare the source objects with objects owned by the target connection user (schema) or with any object accessible by the target connection. The resulting behavior is largely based on which option you select and which objects your connection has access to.

- **Consolidate:** Compares the source objects with those owned by the target schema; in other words the selected objects are "consolidated" into the target schema.

If the target connection does not *own* an object by that name (that is, *target-schema.object-type.object-name*), the object will not be found and a comparison cannot be done. If the target connection does own object *target-schema.object-type.object-name*, *source-schema.objectname* will be compared with *target-schema.object-type.object-name*.

- **Maintain:** Maintains the schema on the source objects and applies it in the target object processing.

If the target connection does not *have access to* those schema objects, they will not be found and a comparison cannot be done. If the target connection does have access to those schema objects, *source-schema.object-type.object-name* will be compared with *target-schema.object-type.object-name*.

Match Constraints By: Controls whether constraints are matched for comparison by the names (**Name**) or definitions (**Definition**).

Ignore Segment Attributes: Causes differences in segment specification to be ignored.

Ignore Table Column Position: Causes differences in the positions of table columns to be ignored when objects are compared. For example, if this option is enabled (checked), two tables would not be considered different if the only difference is that Column1 and Column2 appear first and second in one table but second and first in the other table.

Ignore Storage: Causes differences in storage specification to be ignored.

Ignore High Values: Causes differences in partition high values for tables to be ignored.

Ignore Tablespace: Causes differences in the tablespace specified for the object to be ignored.

Logical Compare: Causes meaningless SQL formatting differences in definitions of objects (such as packages, package bodies, procedures, and functions) to be ignored.

Ignore Partitioning: Causes differences in partitioning specifications for the object to be ignored.

Report Constraint Name Difference: If Match constraints by Definition is selected, causes the constraint name to be displayed when differences are reported.

6.88.2 Types to Diff page

See [Types](#) (page 6-137) under [Shared Wizard Pages](#) (page 6-137).

6.88.3 Specify Objects page

See [Specify Objects](#) (page 6-137) under [Shared Wizard Pages](#).

6.88.4 Summary page

See [Summary](#) (page 6-138) under [Shared Wizard Pages](#) (page 6-137).

The results are displayed in a *Diff Report* window, where you can see the DDL statements to update the objects in the destination schema to reflect differences from the source schema. To create a file containing these DDL statements, click the **Generate Script (SQL)** icon in that window. In the Generate SQL dialog box you can specify the Save As destination or type (Single File, Separate Files, Type Files, Separate Directories, Worksheet, Clipboard), the location, and the encoding.

To display only objects with differences that have the same name in both connections, select **Show Equal Objects**. To display all objects that have differences or that exist in one connection but not the other, select **Show New Objects**.

Related Topics

- [Database Connections](#) (page 1-28)
- [Database Objects](#) (page 1-13)

6.89 Database Export (Unload Database Objects and Data)

Use this wizard to export database objects and optionally data. For a selected database connection, you can export some or all objects of one or more types of database objects to output files, a SQL Worksheet, or the clipboard. The output may contain SQL data definition language (DDL) statements to create these objects. It may also contain SQL statements to insert data into these objects and other formatted files.

Note:

For some needs you can use special SQL Developer worksheet "hints" instead of this wizard to generate output in desired formats, as explained in [SQL Worksheet "Hints" for Formatting Output](#) (page 1-48).

The number of pages and the options available depend on the potential scope of the export operation. For example, if you clicked Tools, then Database Export, all pages are available; however, if you right-clicked a table name and selected Export, fewer pages are required.

In several pages, if you select **Proceed to summary**, clicking Next takes you to the Export Summary page.

See Also:

[Cart Versus Database Export Wizard](#) (page 1-71) to consider whether you may want to use the Cart instead of the Database Export wizard

6.89.1 Source/Destination page

Contains up to four main areas for specifying the database connection and the DDL, data, and target options. If you specify any options, they override the defaults for [Database](#) (page 1-99): Utilities: Export user preferences.

Connection: Select the database connection with the objects to be exported.

Other options on this page: See the [Database](#) (page 1-99): Utilities: Export preferences for explanations of individual options.

Export DDL: If this option is enabled, DDL statements are included in the export operation. Select the options to apply to the DDL that is generated.

Export Data: If this option is enabled, data is included in the operation. Select the options to apply to the data that is unloaded.

Note:

An **Export Format Error** is displayed if you attempt to export both DDL and Data when Format is loader.

Format: Select the desired output format for the data to be unloaded. Depending on the selected format, other options may appear. For example, for **xls** (Microsoft Excel file), you can specify worksheet names for the data and the SELECT statement.

For CLOB data, exporting is supported only if the format is `loader` (SQL*Loader) or `pdf` (PDF). Some export types export only a subset of the string followed by an ellipsis (...).

To *paste unloaded data into a Microsoft Excel* file, specify Export Data but not Export DDL, select **text** for Format, and select **Clipboard** for Save As; and after completing the unload, paste from the clipboard into Excel.

For exporting *date fields in date format* instead of timestamp format: If you encounter problems exporting dates, check that the following line is in the `sqldeveloper\sqldeveloper\bin\sqldeveloper.conf` file to have dates unloaded in date format: `AddVMOption -Doracle.jdbc.mapDateToTimestamp=false`

For exporting *large tables to Microsoft Excel* files:

- If you encounter problems exporting to an `.xls` file, try exporting instead to XLSX (`.xlsx` file) format. However, if you must export to an `.xls` file, try adding the following line to the `sqldeveloper\sqldeveloper\bin\sqldeveloper.conf` file to increase heap size and then restarting SQL Developer: `AddVMOption -Xmx1024M`
- If the number of table rows exceeds 65,536, SQL Developer writes the rows to multiple worksheets within the spreadsheet file.

Save As: Specifies how or where target statements and data are to be saved:

- **Single File:** A single file contains both DDL and data. When you are unloading DDL, only Insert format can be specified for data.
- **Separate Files:** Each object is saved to a separate file in the specified directory.
- **Type Files:** Objects of the same type are saved to a single file in the specified directory.
- **Separate Directories:** A directory for each object type being exported is created. Files are created in the appropriate directory.
- **Worksheet:** Statements are sent to a SQL Worksheet window.
- **Clipboard:** Statements are copied to the clipboard.

6.89.2 Types to Export page

Specify object types to be exported and options for the export operation.

Object Types: Check the types of objects that you want to export. You can click *Toggle All* to check and uncheck all individual types. You must select at least one object type.

6.89.3 Specify Objects page

See [Specify Objects](#) (page 6-137) under [Shared Wizard Pages](#) (page 6-137).

6.89.4 Specify Data page

See [Data](#) (page 6-137) under Shared Wizard Pages.

6.89.5 Export Summary page

See [Summary](#) (page 6-138) under Shared Wizard Pages.

If a script is created and if Save As was specified as Worksheet, the file is displayed in a SQL Worksheet window, where you can run it as a script and perform other operations.

6.90 DDL Panel for Creating or Editing an Object

You can review and save the SQL statement that SQL Developer will use to create or edit the object, to reflect any changes you have made to the object's properties. If you want to make any changes, go back to the relevant panels and make the changes there.

To save the SQL statement to a script file, click **Save** and specify the location and file name.

6.91 Debugger - Attach to JPDA

This dialog box is displayed when you right-click a database connection name and select Remote Debug. Use this dialog box if you are using the Java Platform Debugger Architecture (JPDA) and you would like the debugger to listen so that a debuggee can attach to the debugger.

Host: Name or IP address of the remote host on which SQL Developer should listen for the database to connect.

Port: Listening port number on the remote host. You can choose any valid port number that is not in use by another process.

Timeout: The number of seconds that SQL Developer will wait for the remote database to make a debugging connection.

Don't Show Dialog Box Before Connecting: If this option is checked, this dialog box will not be displayed before future connections for remote debugging.

Related Topics:

[Remote Debugging](#) (page 1-39)

6.92 Deploy or Import Application

Use this wizard to deploy or import an Application Express application into a specified target schema.

Deploy to Connection or Specify File to Import

Choose Connection to Deploy Application: For a deploy operation, specify the database connection for the target schema into which to deploy the application.

Specify File to Import: For an import operation, specify the location and name of the SQL file containing the exported application (usually the output of a previous "export application" operation).

Choose Import Options

Specifies options for the application to be deployed or imported.

Workspace: Name of the Application Express workspace.

Parse As Schema: Schema against which all of the application's SQL and PL/SQL will be parsed.

Application Name: Name of the application.

Application Alias: Alias for the application. It is recommended that you never hard code the application ID into your application, but instead use the application alias or a built-in substitution string (such as APP_ID and APP_ALIAS).

Build Status: RUN_ONLY or RUN_AND_BUILD

Application ID: Specify whether to have an application ID assigned automatically, to use an existing listed ID, or to specify a new ID. Use these options to avoid application ID conflicts, such as when you need to have two versions of the same application in the same instance. For example, you might be migrating an application to a production instance and still need to maintain the development version.

ID Currently Used by and Overwrite: If the specified Application ID is currently used by another application, you can enable **Overwrite** to have the application ID instead associated with the application being deployed or imported.

Summary

Displays the selected options for the application to be deployed or imported. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [Applications \(Application Express 3.0.1 and Later\)](#) (page 1-14)

6.93 Deploy Objects to Cloud

This dialog box is displayed when you click the Deploy Cloud icon in the Cart window.

You can deploy objects to the Cloud, except as noted in [Unsupported Cart Object for Cloud Deployment](#) (page 6-81).

Related Topics:

[Using the Cart](#) (page 1-68)

6.93.1 Cloud

Connection: Name of a previously created Cloud connection. (When you create a Cloud connection, you can use it immediately within the SQL Developer graphical interface; however, to use that connection with the SQL Developer command line interface, you must first exit the SQL Developer session in which you created the connection.)

Title: A short descriptive string identifying this deployment.

Server: Server for the Cloud service (for example, a WebLogic Server system).

Port: Port to be used for the deployment. For example: 2222

6.93.2 Authorization

SFTP User: Secure FTP user name that you received in email when you signed up for the Cloud service.

SFTP Password: Password for the SFTP user.

6.93.3 Transfer

File: .zip file to contain the files to be used by SQL Developer for the deployment.

6.93.4 Deploy DDL

Specify whether not to replace or to replace existing destination objects that have the same schema name and object name combination.

6.93.5 Deploy Data

Truncate destination data: If you specified not to replace existing destination objects, then for each existing destination table, this controls whether any existing data is truncated (deleted) before the source table data is appended.

Disable constraints before moving data: Controls whether to disable referential integrity constraints before deploying data to destination tables. (You may want to disable constraints to expedite the deployment, and then reapply the constraints after the deployment.)

Format

Delimiter: Identifies the character used to separate the data into columns. The delimiter character is not included in the data loaded. If the preview page does not show the data separated into columns, the correct delimiter is probably not specified. Examine the data in the preview area to determine the correct delimiter.

Line Terminator: Identifies the terminator for each line. The line terminator is not included in the data loaded. If the preview page shows the data in one single row, the correct terminator is not specified.

Left Enclosure and Right Enclosure: Enclosures are used for character data and are optional. Enclosures are not included in the data loaded.

6.93.6 Unsupported Cart Object for Cloud Deployment

An error box is displayed if any of the following conditions apply on an attempt to deploy objects from the Cart to the Cloud:

- The objects include any of the following object types: database link, directory, external table, materialized view, materialized view log, or synonym.
- The objects include any with a data type that is not supported for Cloud deployment (such as SDO_GEOMETRY).
- The objects are from a non-Oracle database connection.

6.93.7 Other Errors (Invalid or Missing Entries)

An error box is displayed if any entries are invalid or any required entries are omitted. Check the error text and the dialog box entries, ensure that all required information is entered and all entries are valid, and click **Apply** again.

6.94 Describe Object Window

This window is displayed when you select a database object name in the SQL Worksheet, right-click, and select **Describe**. The information is read-only, and is displayed using tabs that are appropriate for the type of object.

For example, if the display is for a table, the information displayed is similar to that in the [Create/Edit Table \(with advanced options\)](#) (page 6-42) dialog box.

6.95 Diff Objects

This dialog box is displayed if you click the Diff Objects icon in the Cart window.

Open Configuration icon: Opens a previously saved XML configuration file, to use its settings as defaults for this use of the dialog box.

Save Configuration icon: Saves the current settings in the dialog box to an XML file, which you can later open to use for Cart operations of this type.

Destination: Destination with objects to be compared with the source objects: **Cart** for another currently open cart tab, or **Connection** for a database connection that has access to the destination objects to be compared.

DDL Generation Options: Check the options or features to be included if you later generate a script to update the objects in the destination schema to reflect differences in the source schema: comments, schema name (prefixed to schema objects), constraints, password values, constraint indexes, referential constraints. Options that are not checked are ignored or not performed during the script generation.

Schema: This option, if you also select Schema: Consolidate under DDL Comparison Options, causes object names to be prefixed with the target schema name in the generated script (for example, CREATE TABLE "HR"."SKILLS" instead of CREATE TABLE "SKILLS"). If you select Schema: Maintain under DDL Comparison Options, this Schema option is ignored and object names are prefixed with the *source* schema name in the generated script.

DDL Comparison Options: Check various options, including differences to be ignored, when objects in the source and destination schemas are compared.

Note:

The DDL Generation Options apply to how the DDL will be generated, whereas the DDL Comparison Options apply to how the comparison is made between the two objects.

Schema: Consolidate and **Schema: Maintain** enable you either to compare the source objects with objects owned by the target connection user (schema) or with any object accessible by the target connection. The resulting behavior is largely based on which option you select and which objects your connection has access to.

- **Consolidate:** Compares the source objects with those owned by the target schema; in other words the selected objects are "consolidated" into the target schema.

If the target connection does not *own* an object by that name (that is, *target-schema.object-type.object-name*), the object will not be found and a comparison cannot be done. If the target connection does own object *target-schema.object-type.object-name*, *source-schema.objectname* will be compared with *target-schema.object-type.object-name*.

- **Maintain:** Maintains the schema on the source objects and applies it in the target object processing.

If the target connection does not *have access to* those schema objects, they will not be found and a comparison cannot be done. If the target connection does have access to those schema objects, *source-schema.object-type.object-name* will be compared with *target-schema.object-type.object-name*.

Match Constraints By: Controls whether constraints are matched for comparison by the names (**Name**) or definitions (**Definition**).

Ignore Segment Attributes: Causes differences in segment specification to be ignored.

Ignore Table Column Position: Causes differences in the positions of table columns to be ignored when objects are compared. For example, if this option is enabled (checked), two tables would not be considered different if the only difference is that Column1 and Column2 appear first and second in one table but second and first in the other table.

Ignore Storage: Causes differences in storage specification to be ignored.

Ignore High Values: Causes differences in partition high values for tables to be ignored.

Ignore Tablespace: Causes differences in the tablespace specified for the object to be ignored.

Logical Compare: Causes meaningless SQL formatting differences in definitions of objects (such as packages, package bodies, procedures, and functions) to be ignored.

Ignore Partitioning: Causes differences in partitioning specifications for the object to be ignored.

The results are displayed in a *Diff Report* window, where you can see the DDL statements to update the objects in the destination schema to reflect differences from the source schema. To create a file containing these DDL statements, click the **Generate Script** (SQL) icon in that window. In the Generate SQL dialog box you can specify the Save As destination or type (Single File, Separate Files, Type Files, Separate Directories, Worksheet, Clipboard), the location, and the encoding.

To display only objects with differences that have the same name in both connections, select **Show Equal Objects**. To display all objects that have differences or that exist in one connection but not the other, select **Show New Objects**.

Related Topics

- [Using the Cart](#) (page 1-68)
- [Database Connections](#) (page 1-28)
- [Database Objects](#) (page 1-13)

6.96 Drop Pluggable Database

Use this dialog box to drop a PDB.

Database Name: Name of the PDB.

Datafile Action: Determines whether to retain or delete the data files associated with the PDB after the PDB is dropped.

- **Keep:** Retains the data files associated with the PDB after the PDB is dropped. The temp file for the PDB is deleted because it is no longer needed. Keeping data files may be useful in scenarios where a PDB that is unplugged from one CDB is plugged into another CDB, with both CDBs sharing storage devices.
- **Including:** Deletes the data files associated with the PDB being dropped. The temp file for the PDB is also deleted. (Including must be specified if the PDB was created with the SNAPSHOT COPY clause)

Related Topics

[Multitenant Container Database \(CDB\)](#) (page 1-19)

6.97 Drop REST Services for Schema

This dialog box is displayed in you right-click a connection name in the Connections navigator and select Drop REST Services. (This option applies only if Oracle REST Data Services has been installed on the system with the database containing the schema associated with the connection.)

Dropping REST services causes REST access to the schema, including all tables and views in that schema, to be disabled and REST metadata to be deleted.

You are asked if you are sure you want to perform the operation. To perform the operation, click **Apply**; otherwise, click **Cancel**.

6.98 Edit/View Value (Data or Table Column)

This dialog box enables you to view and (if applicable) edit data in certain dialog box fields, in a cell in the table Data grid, in some other grid display (for example, to edit the value of a single column within a row), or in a filter specification. If you are permitted to modify the data, you can change the data value and then click **OK**. (If you are not permitted to modify the data, the Value display is read-only.)

The specific options for editing the value available depend on the context, such as the data type of the column associated with that cell in the grid. For example, for **BLOB** or **BFILE** data you can:

- Load a file containing data (such as an image file)
- Use the external editor associated with the MIME type and file extension

Save the result of any edits

For **date**-based values (such as for **CREATED** as the filter criterion), you can specify a date, time, and time zone, or click **Now** to set the values in the dialog box to the current (as of when you click the Now button) date and time.

Related Topics:

[External Editor](#) (page 1-110)

6.99 Enter Bind Values

This dialog box enables you to enter values for each bind variable. If the NULL option is checked, you cannot enter a value in this dialog box.

6.100 Erase from Disk

This dialog box asks you to confirm your action if you select certain objects (such as one or more files in the Files navigator, or the Connections node in the Connections navigator) and press the Delete key. To perform the deletion, click **Yes**; to cancel the deletion, click **No**.

6.101 Error Writing to Export File

This box is displayed if you tried to export table data to a file, but the directory or folder path does not exist.

Click **OK** to close the box, then enter a valid path in the Export dialog box and click **Apply**.

6.102 Export Error

This dialog box is displayed when you tried to export some or all objects of one or more types of database objects to a file containing SQL statements, but did not include some essential information, which might include one or more of the following:

- The database connection. For **Connection**, select the database connection from which the objects will be exported.
- The name of the output file. Look at the **Options** tab, and be sure that you specified a file.
- One or more objects or types of objects. Look at the **Objects** tab, and be sure that you selected (checked) at least one object or type of object.

6.103 Export Data

This dialog box prompts you to specify the location and name of a text file to contain the output of the export operation, such as data values during a debug operation.

6.104 Export Objects

This dialog box is displayed when you click the Export icon in the Cart window.

Open Configuration icon: Opens a previously saved XML configuration file, to use its settings as defaults for this use of the dialog box.

Save Configuration icon: Saves the current settings in the dialog box to an XML file, which you can later open to use for Cart operations of this type.

Export DDL: If this option is checked, the data definition language (DDL) statements for the database objects to be deployed are included in the output file, and the other options in this group affect the content and format of the DDL statements.

For explanations of specific options, see the Export DDL options under Database: Utilities: Export.

Export Data: If this option is checked, the output file or files contain appropriate statements or data for inserting the data for deployed tables and views. If this option is not checked, the data for deployed tables and views is not included; that is, only the DDL statements are included.

Save As: Specifies how or where the statements and data are to be saved:

- **Single File:** A single file contains both DDL and data. When you are unloading DDL, only Insert format can be specified for data.
- **Separate Files:** Each object is saved to a separate file in the specified directory.
- **Type Files:** Objects of the same type are saved to a single file in the specified directory.
- **Separate Directories:** A directory for each object type being exported is created. Files are created in the appropriate directory.
- **Worksheet:** Statements are sent to a SQL Worksheet window.
- **Clipboard:** Statements are copied to the clipboard.

Compressed: If this option is checked, data segments are compressed to reduce disk use.

Encoding: Character set for the statements and data.

File: For exports to one or more files, the .zip file to contain the statements to perform the operation.

Related Topics:

[Using the Cart](#) (page 1-68)

[Save or Open Cart Tool Configuration](#) (page 6-134)

[Database](#) (page 1-99)

6.105 Export Resource Module, Privilege, and Role

In the Connections navigator, using REST Data Services, you can export the resource module, privilege, or role definitions for a schema associated with a connection. The exported definitions are saved to a SQL file. Based on the options selected in the Export dialog, the exported file contains the definitions for Enable Schema, the resource module, the resource template, the resource handlers, the privileges associated with the module, and the roles associated with the privileges.

The Export dialog is displayed when you right-click the module name, privilege name, or role name in the Connections navigator. If you want to export all the modules, privileges, or roles, right-click the Modules, Privileges, or Roles node.

- **Enable Schema:** Exports the definition for Enable Schema. Enable Schema controls whether Oracle REST Data Services access is enabled for the specified schema or object.
- **Privileges (includes protected module, roles, and resource patterns):** Exports the privilege definitions associated with the module.

- **Save:** Specifies the location to save the SQL file.
- **File Encoding:** Specifies the character encoding for the file.

Related Topics

- [REST Data Services in the Connections Navigator](#) (page 3-14)
- [RESTful Services Wizard and Object Dialog Boxes](#) (page 6-127)

6.106 Export: Unique Names Warning

If a cart to be exported contains multiple objects with the same name and different owners (schemas), and if the Show Schema option is not enabled, you are warned that some objects in the cart are not unique and that all objects must be uniquely identified by object type and name. (The Show Schema option causes the schema name to be included in CREATE statements: for example, CREATE TABLE SMITH.BOOKS instead of CREATE TABLE BOOKS.)

For example, if the cart contains the tables SMITH.BOOKS and JONES.BOOKS, and if the Show Schema option is not enabled, you cannot include both tables in the export operation.

To perform the export operation, either enable the Show Schema option or select only one of the objects with the same name. If you click **Yes** to continue, the Show Schema option is selected in the [Export Objects](#) (page 6-85) dialog box and cannot be deselected.

6.107 External Locator Configuration

This dialog box is displayed if you click External Locator Configuration when creating a repository. Specify the information required to connect to the remote repository when the method by which the client will gain access to and authenticate against the server is External.

Set Remote Shell: If this option is checked, external repositories are accessed through a remote shell utility, usually rsh (the default) or ssh.

Set Remote Shell: If this option is checked, you specify the name of the program on the remote server. (It is unlikely to need to be changed from the default, and should only be changed in cooperation with the administrator of the remote server.)

Related Topics

- [Using Versioning](#) (page 1-53)

6.108 External Tools

This dialog box is displayed when you click Tools and then External Tools. It displays information about user-defined external tools that are integrated with the SQL Developer interface.

Find Tools: Checks for any tools that Oracle offers for your consideration, and adds them to the list if they are not already included.

New: Starts a wizard for defining a new external tool.

Edit: Displays a dialog box for editing the selected external tool.

Related Topics:

[Create/Edit External Tool](#) (page 6-88)

6.109 Create/Edit External Tool

This interface is displayed as a wizard if you are creating a new external tool, and as a dialog box if you are editing an existing external tool.

Related Topics:

[External Tools](#) (page 6-87)

6.109.1 Program Options

Program Executable: Path of the program executable for the tool.

Arguments: Arguments (parameters) to be passed to the program. You can click **Insert** to insert a macro for the argument.

Run Directory: Directory in which to run the program. You can click **Insert** to insert a macro for the directory.

Command Sample: A read-only sample display of the command to run the program.

Related Topics:

[Insert Macro](#) (page 6-100)

6.109.2 Display

Specify how the external tool should appear when displayed in menu or toolbar items.

Caption for Menu Items: The text string that will appear for any menu item that calls the external tool. To indicate the mnemonic character, use the ampersand before the character. For example: &Mytool for the "M" to be underlined and used as the mnemonic.

ToolTip Text: Text for the tooltip to be displayed when the mouse pointer hovers over the icon for the tool in a toolbar.

Icon Location: File path of the icon associated with the tool. Click Browse to specify a graphics file, or Use Default to use the default icon (if you previously specified a nondefault icon).

Preview: A read-only display of the menu item and its associated icon.

6.109.3 Integration

Specify how the external tool will be integrated with SQL Developer.

Add Items to Menus: Check any menus on which you want to include an item for this tool.

Add Buttons to Toolbars: To add the icon for this tool to the SQL Developer main toolbar, check Main Toolbar.

After Tool Exits: To have SQL Developer reload any open files after the tool exits, check Reload Open Files.

6.109.4 Availability

Specify when the external tool is enabled. In contexts where the tool is not enabled, its menu item and icon are grayed out.

Always: Makes the tool always available.

When a File is Selected or Open in the Editor: Makes the tool available only when a file is selected or open, such as when the SQL Worksheet is open.

When Specific File Types are Selected: Makes the tool available only when files of the specified type or types are selected. Use the arrow buttons to move desired types from Available Types to Selected Types.

6.110 Choose Offline Options

This dialog box is displayed when you click Tools, then Migration, then Create Database Capture Scripts. It specifies options for creating an offline capture properties (.ocp) file, which you can later specify as the Offline Capture Source File on the Source Database page of the Migration Wizard.

Output Directory: Converted model containing tables whose data is to be moved to the corresponding Oracle database tables.

Generate for: *Windows Batch File* generates a .bat file to be run on Windows systems; *Linux Shell Scripts* generates .sh files to be run on Linux systems.

For a MySQL migration, if you generate .sh files, you must also execute the following command to make the .sh files executable and the .ocp file writable:

```
chmod 755 *
```

Platform: The IDM DB2, MySQL, Microsoft SQL Server, Sybase Adaptive Server, or Teradata version for which to generate the scripts.

Related Topics:

[Source Database](#) (page 2-3)

[SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

6.111 Edit Font Location

Lets you specify the locations for fonts to be used in PDF generation, and the order in which the locations are searched.

Add Entry: Lets you select a location to be added to the list.

Edit Entry: Lets you change the location for the selected item in the list.

Remove: Deletes the selected location from the list.

Up and down arrows: Let you move selected items up or down to change the order of items in the list.

Related Topics

- [Database](#) (page 1-99): Utilities: Export: Formats (CSV, Delimited, Excel, Fixed, HTML, PDF, SQL*Loader, Text, XML) user preferences

6.112 Edit Header or Footer (PDF Formatting)

In the user preferences for **PDF Format: Header and Footer** under Database: Utilities: Export: Formats (CSV, Delimited, Excel, Fixed, HTML, PDF, SQL*Loader, Text, XML), you can edit the header and footer separately to specify text to be used.

To include the date in that text, click the Insert Date icon to display a dialog box to specify the desired date format, and whether or not to have the inserted date be updated automatically to reflect any update of the PDF file.

Related Topics:

[Database](#) (page 1-99)

6.113 Edit Join

This dialog box lets you edit the join specification for a join view.

Swap: Reverses the order of the tables.

Join Type:

Natural: If this option is checked, a natural join will be performed. A natural join is based on all columns in the two tables that have the same name. It selects rows from the two tables that have equal values in the relevant columns. When specifying columns that are involved in the natural join, do not qualify the column name with a table name or table alias.

On: Using the ON clause to specify a join condition lets you specify join conditions separate from any search or filter conditions in the WHERE clause.

Using: When you are specifying an equijoin of columns that have the same name in both tables, the USING *column* clause indicates the columns to be used. You can use this clause only if the join columns in both tables have the same name. Within this clause, do not qualify the column name with a table name or table alias. In an outer join with the USING clause, the query returns a single column which is a coalesce of the two matching columns in the join.

Related Topics

- [Create/Edit View](#) (page 6-62)

6.114 Feature Missing

This dialog box is displayed if you try to use a SQL Developer feature that requires that the specified Oracle Database feature also be installed. To use the SQL Developer feature, install the Oracle Database feature if you are authorized to do so; otherwise, see the database administrator for assistance.

6.115 Feature Required

This dialog box is displayed if you try to use a SQL Developer feature that requires the licensing of the specified feature for Oracle Database. If you do not have a license for the specified feature, you must click **No**.

If you have a license for the feature on the database or databases on which you plan to use the feature this time, you can click **Yes**.

To manage the license information for any database connections (and to minimize prompts for license-related information), you can specify user preferences for [Database](#) (page 1-99): Licensing.

To purchase any required license, contact your Oracle sales representative or authorized Oracle Reseller, or go to the Oracle Store to buy online.

6.116 File Processing

This dialog box is displayed if you right-click a directory or file in the Files navigator and select Tag SQL. In the selected file or in files under the selected directory, all SQL statements that match a specified pattern are rewritten so that a comment is inserted after the SELECT keyword in the main query block that uniquely tags the statement.

For example, a PL/SQL source containing the statement "SELECT 1 FROM dual" would be rewritten into something like "SELECT /* PREFIX 00f7d2 */ FROM dual". This allows easy identification of problematic statements that appear in performance views such as V\$SQL.

Tag Prefix: Text to appear after "PREFIX" in inserted comments.

Selected Extensions: Extensions of files on which to perform the SQL statement rewrite operations.

6.117 Filter

This dialog box is displayed when you right-click a connection node or an object type node (such as Tables) in the Connections navigator and select Apply Filter. Use this box to limit the number of objects of that type that are displayed, according to one or more filter criteria that you specify. For each criterion, specify the following:

- Criterion name (list always includes NAME; other criteria depend on the object type)
- Operator (for example, LIKE)
- Value for comparison (for example EM%)
- Case-sensitive option for character data comparison

For example, to display only tables with names that start with EM, specify NAME, LIKE, and EM% (with the percent sign as a wildcard character).

Any remaining options in this dialog box depend on the context for the filter option.

Override Schema Filter (object type node filters): If this option is checked, any filter criterion specified at the connection level is ignored, and only the object type node filter criteria are applied.

Include Synonyms (object type node filters): If this option is checked, synonyms for objects of this object type are included.

To add another filter criterion, click the Add (+) icon; to delete a criterion, select it and click the Delete (X) icon; to move a criterion up or down in the list, select it and use the arrow icons.

To apply the filter criteria to the Connections navigator display, click **OK**.

To remove the effects of applying a filter, right-click the object type node in the Connections navigator display and select **Clear Filter**.

6.118 Filter Object Types

This dialog box filters (restricts) the types of objects to be displayed for the schema associated with the selected user.

Available Object Types: Lists the types of objects that are available to be added to the display.

Displayed Object Types: Lists the types of objects that are included in the display.

To add a type of object to the display, select it in Available Object Types and click the Add (>) icon; to remove a type of object from the display, select it in Displayed Object Types and click the Remove (<) icon. To move all types of objects from available to displayed (or the reverse), use the Add All (>>) or Remove All (<<) icon.

6.119 Filter Schemas

This dialog box enables you to restrict the schemas that are displayed under Other Users in the Connections navigator.

Available Schemas: Lists the schemas that are not currently displayed under Other Users in the Connections navigator, but that are available to be added to the list of displayed users.

Displayed Schemas: Lists the schemas that are to be included in the display under Other Users in the Connections navigator.

To add a schema to the display, select it in Available Schemas and click the Add (>) icon; to remove a schema from the display, select it in Displayed Schemas and click the Remove (<) icon. To move all schemas from available to displayed (or the reverse), use the Add All (>>) or Remove All (<<) icon.

Only display schemas with visible objects: Limits the display to available schemas that have any database objects that are visible to the database user associated with the current connection.

6.120 Filter Error

This dialog box is displayed if you did not specify any data for an export operation. Be sure to specify Filter Data options that select some data for the export operation.

6.121 Find/Highlight

This dialog box is displayed when you right-click in the table data grid and select Find/Highlight to find, and optionally highlight with a specified background color, occurrences of a specified string value in the table data. Type the string value, and press the down-arrow to search forward from the currently selected grid cell or the up-arrow to search backward.

Search goes to matching cells but does not highlight; **Highlight** goes to matching cells and highlights.

Text matching options: You can ignore case, require a whole word match, require a starting with match, or allow the value to wrap across cell boundaries.

Color (Highlight only): Specifies the background color for highlighting.

Persist Highlight (Highlight only): If this option is enabled, matches are added to the Persisted Highlights list.

Highlight Row (Highlight only): If this option is enabled, the entire row that contains the cell with a matching value is highlighted.

6.122 Git: Add

Adds the selected file to the Git staging area. Once a file is in the staging area, it is ready for you to commit to the repository.

If the file is open (that is, not yet saved), you are prompted to save the file before adding it.

Name and Location: Lists the name and location of the file to be added to the staging area.

Related Topics

- [Using Versioning](#) (page 1-53)

6.123 Git: Add All

Adds all files not yet added to the Git staging area. Once added to the staging area, these files are ready to be committed to the Git repository.

If any files have not yet been saved, you are prompted to save them.

Name and Location: Lists the names and locations of the files to be added to Git.

Related Topics

- [Using Versioning](#) (page 1-53)

6.124 Git: Add to .gitignore File

Lets you mark a file, or a pattern that identifies common file names, as content that Git should ignore.

Often, a directory contains files that should not be kept under version control. For example, log files from a debug or batch operation do not need to be tracked or merged, yet they are often in the same directory as the shared code for a project. Such files should be marked to be ignored by Git.

Related Topics

- [Using Versioning](#) (page 1-53)

6.125 Git: Checkout Revision

Checks out files from a Git repository.

When you check out files from the Git repository, you can check out a specific revision and branch. Optionally, you can select from existing tags and check out a tagged revision. You can also check out to a specific commit; by default, the most recent commit is used at checkout

Name and Location: Name and location of the local repository to which you are checking out files.

Branch: The branch you are using for this checkout. If you know the name, type the branch name. Otherwise, click **Select Branch** to browse available branches. If you are creating a new branch, click **Create Branch**. If you check out a remote branch without

specifying a new local branch to track the changes, the changes will become disassociated from the original branch.

Tag: You can enter a tag to help you select the desired revision for checkout. If you know the name of the tag you want to use, type the tag name. Otherwise, click **Select Tag** to browse the list of available tags

Use Commit ID: The ID for the commit for which this checkout is to be used. Click **Select Commit** to browse from available commit options.

Create Branch: Creates a new branch to use for this checkout.

Related Topics

- [Using Versioning](#) (page 1-53)

6.126 Git: Clone from Git

Clones a Git repository into a newly created directory, creates remote-tracking branches for each branch in the cloned repository, and creates and checks out an initial branch that is forked from the cloned repository's currently active branch.

Based on the information that you enter on these pages, SQL Developer executes the `git clone` command. For detailed information, see the *Git Reference Manual* at <http://www.git-scm.com/documentation> and the `git-clone(1)` manual page at <https://www.kernel.org/pub/software/scm/git/docs/git-clone.html>.

Remote Repository

Specify the information required to log in to the remote Git repository for accessing the files for the operation.

Remote Name: Identifier that you will use when referring to the remote repository. Enter a descriptive, unique name for the clone you are creating.

Repository URL: The URL of the file system at which the Git repository resides. Your Git administrator should be able to provide you with this information.

User Name: If the remote repository does not allow anonymous read access, enter your user name for the server in this field. If your repository permits anonymous access, you can skip the remaining fields on this screen and click **Next** to continue.

Password: If your repository requires a login with user name and password, select this option, then enter your password in the field. You can then click **Next** to continue.

Private Key File: If your repository connection uses a private key file, select this option, then enter the path to the private key file. You can click **Browse** to select the file from a standard directory browser.

Passphrase: If your private key file requires a passphrase, enter the passphrase.

Remote Branch

Specify the branches that you want to include in your clone of the remote repository. Each branch in the repository is represented on the right side. To select a branch, check the box next to the branch name.

Destination

Specify all information for the destination in your local system to which you want to copy your Git repository content.

Destination: Pathname on your local system to which you want the repository to be cloned. You can either type the pathname in the field, or click **Browse** to select the file.

Clone Name: Name for the clone you are creating.

Checkout Branch: Branch to use for the clone you are creating.

Summary

Displays the selected options for the operation. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [Using Versioning](#) (page 1-53)

6.127 Git: Commit

Commits a file (located in the staging area) to your Git repository. If any files have not yet been saved, you are prompted to save the files before this dialog box is displayed.

You can commit an individual file or selected multiple files. If you want to commit all uncommitted files (as opposed to selecting from all the available files) in one operation, select Team > Git > Commit All.

Name and Location: Name and physical locations of the file that will be committed to the repository.

Commit Non-Staged Files: Lets you commit a file that you have not yet staged (that is, files not yet on a staged index list).

Comments: Comments to accompany the commit action. You will later be able to see these comments when viewing the list of versions of a particular file.

Template or Previous Comments: A template with comment text for the Comments box. You can make general changes and additions to the comment templates by clicking the link to comment templates.

Related Topics

- [Using Versioning](#) (page 1-53)

6.128 Git: Commit All

Saves and commits to the Git repository all open and uncommitted files at the same time.

If you want to commit multiple files, but not necessarily all files not yet committed, you can select them (using Shift-click) from the Applications pane, then select Team > Git > Commit.

Name and Location: Name and physical locations of the file that will be committed to the repository.

Commit Non-Staged Files: Lets you commit a file that you have not yet staged (that is, files not yet on a staged index list).

Comments: Comments to accompany the commit action. You will later be able to see these comments when viewing the list of versions of a particular file.

Template or Previous Comments: A template with comment text for the Comments box. You can make general changes and additions to the comment templates by clicking the link to comment templates.

Related Topics

- [Using Versioning](#) (page 1-53)

6.129 Git: Create Branch

Creates a branch, beginning with an existing branch. Note that you can also create a new branch by checking out files from your repository and giving the checked-out files a new branch name.

Name: Name of the branch to be created.

Branch: Branch against which you intend to use this tag. You can type the branch name into the field, or you can click **Select Branch** to browse available branches.

Tag: You can add a tag to this branch when you create the branch. If you know the name of the tag you want to use, type the tag name. Otherwise, click **Select Tag** to browse the list of available tags.

Use Commit ID: The ID for the commit for which this branch will be created. Click **Select Commit** to browse from available commit options.

Related Topics

- [Using Versioning](#) (page 1-53)

6.130 Git: Create Tag

Creates a tag, selecting from available branches and existing tags.

Name: Name of the tag to be created.

Comments: A comment about this tag, such as the release it is intended to support, a bug it is meant to fix, or some other identifying information that will help you select it properly in future.

Branch: Branch against which you intend to use this tag. You can type the branch name into the field, or you can click **Select Branch** to browse available branches.

Tag: You create a tag which is a subset of an existing tag. If you know the name of the tag from which you want to create a subset, type the tag name. Otherwise, click **Select Tag** to browse the list of available tags.

Use Commit ID: The ID for the commit to which this tag will be applied. Click **Select Commit** to browse from available commit options. The default commit ID is the most recent one created.

Related Topics

- [Using Versioning](#) (page 1-53)

6.131 Git: Export Committed Changes

Creates a file containing the changes in all committed files. The file is displayed in the SQL Developer editor. You can specify the following conditions pertaining to the export.

Name and Location: Names and locations of files for which you are exporting changes.

Export File: Name of the file to which the changes will be exported.

Branch: Branch name to use for this export. To select from a list of available branches, click **Select Branch**.

Tag: Click **Select Tag** to choose from a list of available tags to use for this export.

Use Commit ID: The commit ID to use for this export (optional). To select from a list of available commit IDs, click **Browse**.

Related Topics

- [Using Versioning](#) (page 1-53)

6.132 Git: Export Uncommitted Changes

Exports changes that you have not yet committed to a text file, which will be viewed in SQL Developer. You can save this file if you want to share it among team members or make other use of the list of changes.

The files from which these changes are taken must reside in the staging area: that is, they must have been added with the Add or Add All commands, but not yet committed. Any files that you have edited but not yet added to the staging area will be ignored.

Name and Location: Names and locations of files for which you are exporting changes.

Select: Check this box to include the file in the export.

Export File: Name of the file to which the changes will be exported.

Related Topics

- [Using Versioning](#) (page 1-53)

6.133 Git: Fetch from Git

Fetches remote sources into your local Git repository.

Fetching a repository copies changes from the remote repository into your local system, without modifying any of your current branches. Once you have fetched the changes, you can merge them into your branches or simply view them.

Remote Repository

Specify the information for the remote repository from which you want to fetch changes.

Remote Name: Identifier that you will use when referring to the remote repository. Enter a descriptive, unique name.

Repository URL: The URL of the Git repository from which you are fetching files. Your Git administrator should be able to provide you with this information.

User Name: If the remote repository does not allow anonymous read access, enter your user name for the server in this field. If your repository permits anonymous access, you can skip the remaining fields on this screen and click **Next** to continue.

Password: If your repository requires a login with user name and password, select this option, then enter your password in the field. You can then click **Next** to continue.

Private Key File: If your repository connection uses a private key file, select this option, then enter the path to the private key file. You can click **Browse** to select the file from a standard directory browser.

Passphrase: If your private key file requires a passphrase, enter the passphrase.

Remote Branch

Specify the branches to include when you fetch changes from the remote repository. Each branch in the repository is represented on the right side. To select a branch, check the box next to the branch name.

Note that the changes you fetch will not affect any branches in your local repository. You can review the changes, then either ignore them or merge them into your local repository.

Summary

Displays the selected options for the operation. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [Using Versioning](#) (page 1-53)

6.134 Git: Initialize Repository

Creates a Git repository.

Repository Path: Location for the new repository. You can click **Browse** to select the location.

Related Topics

- [Using Versioning](#) (page 1-53)

6.135 Git: Merge

Merges changes, optionally specifying the branch, tag, and commit IDs to use for the merge.

Name and Location: Names and locations of files to be merged.

Branch: Branch you want to merge with the files in your local system. To select from a list of available branches, click **Select Branch**.

Tag: You can sort the target of the merge by tags, if you have set up tags to keep track of your projects. Click **Select Tag** to view a list of available tags.

Use Commit ID: The ID for the commit for which this branch will be merged. To select from a list of available commit IDs, click **Browse**.

Related Topics

- [Using Versioning](#) (page 1-53)

6.136 Git: Pull from Git

Pulls remote sources into your local Git repository.

Pull automatically tries to merge the files you are pulling with any files that are already in your local repository. If you are concerned that this will cause merge conflicts and issues, a safer method is to use Fetch, and then Merge the files into your local repository.

Remote Repository

Specify the remote repository used for pulling source files into your Git repository.

Remote Name: Identifier that you will use when referring to the remote repository. Enter a descriptive, unique name, or select from the list if you have more than one available.

Repository URL: The URL of the file system at which the Git repository resides. Your Git administrator should be able to prove you with this information, or select from the list if you have more than one repository available.

User Name: If the remote repository does not allow anonymous read access, enter your user name for the server in this field. If your repository permits anonymous access, you can skip the remaining fields on this screen and click **Next** to continue.

Password: If your repository requires a login with user name and password, select this option, then enter your password in the field. You can then click **Next** to continue.

Private Key File: If your repository connection uses a private key file, select this option, then enter the path to the private key file. You can click **Browse** to select the file from a standard directory browser.

Passphrase: If your private key file requires a passphrase, enter the passphrase.

Remote Branch

Specify the remote branches to pull into your local repository.

Include: Check the branch you want to pull from.

From: Branch, in the remote repository, from which to pull files.

To: Branches that will be part of the pull operation.

Summary

Displays the selected options for the operation. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [Using Versioning](#) (page 1-53)

6.137 Git: Push to Git

Copies files from your local system to the remote Git repository.

When preparing to push changes, it is recommended that you fetch and merge the latest changes from the remote repository. If there is any merge conflict on the push, the operation will fail.

Remote Repository

Specify the remote repository for the push operation.

Remote Name: Identifier that you will use when referring to the remote repository. Enter a descriptive, unique name, or select from the list if you have more than one available.

Repository URL: The URL of the file system at which the Git repository resides. Your Git administrator should be able to prove you with this information, or select from the list if you have more than one repository available.

User Name: If the remote repository does not allow anonymous read access, enter your user name for the server in this field. If your repository permits anonymous access, you can skip the remaining fields on this screen and click **Next** to continue.

Password: If your repository requires a login with user name and password, select this option, then enter your password in the field. You can then click **Next** to continue.

Private Key File: If your repository connection uses a private key file, select this option, then enter the path to the private key file. You can click **Browse** to select the file from a standard directory browser.

Passphrase: If your private key file requires a passphrase, enter the passphrase.

Local Branch

Specify the branch in your local file system that you will push to the remote repository.

Include: Check all local branches that you want to include in the push operation.

From: Local branch you want to push to the repository.

To: The name that will be used in the remote repository after the push is completed.

Status: The status for this push: `Update` (if the push will result in updating remote files) or `Create` (if the push will result in creating new files in the remote repository).

Summary

Displays the selected options for the operation. To make any changes, click **Back**. To perform the operation, click **Finish**.

Related Topics

- [Using Versioning](#) (page 1-53)

6.138 Git: Revert

Undoes changes made since the last checkout of the file.

Name and **Location:** Lists the names and physical locations of the file that you are about to revert to the last version stored in the repository.

Related Topics

- [Using Versioning](#) (page 1-53)

6.139 Insert Macro

This dialog box is displayed when you click **Insert** when specifying external program options (see [Create/Edit External Tool](#) (page 6-88)). It enables you to insert a sample text string into the relevant field for the external program option; you can then edit that string to suit your needs. (This is somewhat analogous to using snippets to insert text strings into the SQL Worksheet.)

Select the desired type of macro, read its description to ensure that it is what you want, and click **OK**. For some macros, a sample expansion is included.

6.140 Externally Modified Files

This dialog box filters is displayed when an external application has modified a file that you have open in SQL Developer. You are asked if you want to reload the externally modified file.

If you are offered the choice of **Yes** or **No**: If you click **Yes**, the externally modified file overwrites any changes that you might have made in SQL Developer. If you click **No**,

the externally modified file will be overwritten by your version when you save the file in SQL Developer.

If you are offered only the choice of OK (or Help): Click **OK** and then do one of the following:

- To discard changes in the Worksheet and replace the Worksheet with the file currently on disk, click **File**, then **Replace With**, then **File on Disk**.
- To replace the file on disk with the contents of the Worksheet, click **File**, then **Save**.

6.141 Find/Replace Text

This dialog box specifies a text string to find, optionally a replacement text string, and search options.

Text to Search For: Text string to search for.

Replace With: If you check this option, enter a text string to replace the text string that is being searched for.

Options: Options to control the search behavior: **Match Case** makes the search case-sensitive; **Search from Beginning** starts the search at the beginning instead of at the text cursor; **Highlight All Occurrences** highlights all occurrences of the search string instead of just the first one; **Wrap Around** searches across line breaks; **Whole Word Only** find the search string only if it is a complete word and not just part of a word; **Regular Expressions** means that the search string is a regular expression; **Selected Text Only** means to search only in the text block that you have selected.

Direction: **Forward** starts the search from the cursor in the direction of normal text flow; **Backward** starts the search from the cursor in the opposite direction of normal text flow.

6.142 Find Result

This box is displayed if you specify text to search for in the [Find/Replace Text](#) (page 6-101) dialog box that is not in the SQL Worksheet.

If you think that the text is in the worksheet, retry your query, and check the spelling of the text to search for.

6.143 Format Properties (Advanced Formatting)

The Format Properties box is displayed if you right-click and select Advanced Format in the SQL Worksheet or the editor for a subprogram, package, view, or trigger. You can specify a set of formatting rules different from those if you had selected Format (that is, different from the [Database](#) (page 1-99): SQL Formatter user preferences).

Output Destination: **Editor** applies the selected Output Type formatting in the current editing pane; **Clipboard** applies the formatting on the clipboard, so that you can paste it into a pane or window of your choice.

Output Type: A set of formatting rules associated with an output type: SQL or another type from the list. You can test several output types on a statement in the SQL Worksheet (pasting or reentering the statement each time). For example, if the original statement is entered as:

```
select 1 from dual
```

The result with the following output types is:

- SQL: `SELECT 1 FROM dual`
- Java StringBuffer: `SQL.append("SELECT 1 FROM dual");`

Enclosed In: Not Enclosed (for example, `SELECT 1 FROM dual`), " (for example, when embedded in Java: `SQL.append("SELECT 1 FROM dual");`) or, ' (for example, in MSSQL string: `N'SELECT 1 FROM dual'`)).

Escaped By: A character that can be used for "escaping" the Enclosed In character. When Enclosed In is " or ', that character can be escaped by itself or by the Escaped By character. For example, if Escaped By is \ and the input is `SQL.append("SELECT 1 \"one\" FROM dual");`, it will be formatted as `SELECT 1 "one" FROM dual`.

6.144 Generate Patch

You can generate a patch containing changes that have been made to files. The patch can be applied to another set of checked out files so that your changes are incorporated into them. A patch must be applied to the same revision/tag from which it was generated.

Source Files or Project: Name and location of the project or set of files will be the subject of the patch.

Patch Target: Where the generated patch will be sent: **System Clipboard** (from which you can paste the patch), **File** (accept the shown name and location, or specify different ones), or **Open Patch File Editor** (where you can edit and then save it).

Output Format: One of the standard diff formats: **Unified**, **Context**, or **Standard**.

6.145 Go to Bookmark

Use this box to specify the bookmark to go to in the selected function or procedure. After you enter the bookmark and click **Go**, the line associated with that bookmark is highlighted.

6.146 Go to Line Number

Use this box to specify the line number to go to in the selected function or procedure. After you enter the line number and click the Go icon, that line is highlighted.

6.147 Go to Line Number: Error

This error box tells you that you entered an invalid line number in the Go to Line Number box, probably because you entered a line number greater than that of the last line in the function or procedure.

6.148 History

The History window displays the differences between different revisions of the same file. You can see the local history of a file and, when appropriate, the version history of a file.

Revision Filter: Use to filter the list of revisions displayed. You can edit any of the filters, or add new filters, using the Customize Filters button.

Customize Filters button: Opens the Custom Filters dialog box, in which you can edit the filter values of existing filters or add new filters.

Refresh button: Updates the display.

Go to First / Previous / Next / Last Difference buttons: Moves the cursor to the first, previous, next, or last difference.

Purge Local History button: Removes the local history versions of the files.

Restore from Revision button: Makes the current version of the file (in the right-hand area) the same as the revision of the file (in the left-hand area).

Generate Patch button: Opens the Generate Patch dialog, where you can generate a patch containing changes that have been made to the files.

Left-hand and right-hand boxes: The versions are aligned line by line. Lines with differences are highlighted using shaded boxes, joined as appropriate. You cannot edit the left-hand area; you can edit the right-hand area.

Related Topics:

[Custom Filters](#) (page 6-70)

[Generate Patch](#) (page 6-102)

[Using the SQL Worksheet](#) (page 1-41)

6.149 Import Using Oracle Loader for Hadoop

Use this wizard to initiate an Oracle Loader for Hadoop job. To use the wizard, you must understand the information in [Apache Hadoop Connectors Support in SQL Developer](#) (page 1-79), and especially [Oracle Loader for Hadoop: Overview and Requirements](#) (page 1-81).

Source

Specify whether you are loading from delimited text files or a Hive table. If you are loading from a Hive table you can select the table to load from by selecting the

- Hive connection
- Hive schema
- Hive table

You must select an SSH connection. This SSH connection can be to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. The Oracle Loader for Hadoop job will be run from this node.

Text Preview (if Source Type is Text)

You can select the HDFS input directory and data files to load. Data files in one directory can be loaded in one job. For loading data from multiple directories, use Oracle SQL Connector for HDFS from the command line or use Oracle Shell for Hadoop Loaders.

You can also select the file format, delimiter, and other related parameters.

The input data columns in the text file are listed in order from left to right, with the labels F0, F1, ..., Fn.

Column Properties

You can map the source data columns to database table columns.

Pay attention to the data format (a common source of errors). In particular, pay attention to the proposed date format mask. The date mask that appears by default is a “best guess” only.

Ensure that any conflicts and warnings are resolved before you go to the next page.

Where

You can select the data you want to copy. You can filter rows and partitions with a WHERE clause. If data is loaded from a partitioned Hive table, the partition keys and partitions are listed for your reference.

After you make your selections, click **Refresh** to see sample rows that match your selection.

Additional Properties

You can select

- **Import method:** If the database table is partitioned, loading with Direct Path method is recommended for best performance. If the database table is not partitioned, only the JDBC method can be used.
- **Parallelism:** This value should be chosen depending on the number of reduce tasks available in the Hadoop cluster. The higher the value, the faster the load will be.

These two specifications determine the performance of the load.

HDFS Job Directory: The HDFS directory location for storing log files, .bad files with rejected rows, and the Oracle Loader for Hadoop job report. This directory is created by SQL Developer when running Oracle Loader for Hadoop.

Destination Connection: If you are using Oracle Wallet, specify credentials. Otherwise, the credentials from the SQL Developer connection will be used.

Summary

You can review the Oracle Loader for Hadoop script by clicking the Script tab before clicking **Finish** to submit the job.

6.150 Import Using Oracle SQL Connector for HDFS

Use this wizard to load data into an Oracle Database table using SQL and external tables, and to query data in the Hadoop Distributed File System (HDFS) in-place using external tables. To use the wizard, you must understand the information in [Apache Hadoop Connectors Support in SQL Developer](#) (page 1-79), and especially [Oracle SQL Connector for HDFS: Overview and Requirements](#) (page 1-82).

Source

Source Type: Specify whether you are loading from delimited text files or a Hive table. Note that the Hive table should be over text files.

If you are loading from a Hive table you can select the table to load from by specifying the **Connection**, **Schema**, and **Table**.

SSH Host: You must select an SSH connection. This SSH connection can be to a Hadoop client, a Hadoop edge node, or a node in the Hadoop cluster. The Oracle SQL Connector for HDFS job will be run from this node.

Text Preview (if Source Type is Text)

You can select the HDFS input directory and data files to load. Data files in one directory can be loaded in one job. For loading data from multiple directories, use Oracle SQL Connector for HDFS from the command line or use Oracle Shell for Hadoop Loaders.

You can also select the file format, delimiter, and other related parameters.

The input data columns in the text file are listed in order from left to right, with the labels F0, F1,...,Fn.

Import Method

Import Method: for **partitioned** Hive tables:

- **External Table** (for query): A metadata table (Metadata Table), an external table for each partition, and a view for each external table are created. The metadata table lists the external table name and view name for each Hive table partition. Select the name of the Metadata Table; the names of the external tables and views are derived from this name.
- **Staging External Table** (for load): A metadata table (Metadata Table), an external table for each partition, a view for each external table, and a database table (Destination Table) are created. You can drop the metadata table, external tables, and views after completion of the load. Select the name of the Destination Table; all the other names are derived from this name.

Import Method: for **non-partitioned** Hive tables:

- **External Table** (for query): An external table (Destination External Table) is created. This can be used to query the data in-place in HDFS.
- **Staging External Table** (for load): An external table (Staging External Table) and a database table (Destination Table) are created, with the option of dropping the Staging External Table after completion of the load. The data is accessed through the external table and loaded into the database table.

Column Properties

You can modify the definitions of columns in the destination table. When the input source is Hive, the default mappings are selected from the Hive table definition.

Source Data Columns, Target Data Columns, Data Type: Specify Oracle data type definitions if you would like to change any of the default mappings. For example, Hive date columns are often of type 'string' and will be mapped to VARCHAR2. You can change that mapping to DATE or TIMESTAMP depending on your data.

Ensure that any conflicts and warnings are resolved before you move to the next page.

Where (for partitioned Hive tables)

You can filter rows with a WHERE clause. You can only filter rows when the input Hive table is partitioned, and you can only use the partitioning key in the filter clause.

Additional Properties

Preprocessor Directory: The database directory that points to the installation location of \$OSCH_HOME/bin on the database system. This directory is created when Oracle SQL Connector for HDFS is installed on the database system. The directory name is typically OSCH_BIN_PATH.

Parallelism: Determines the performance of load or query using Oracle SQL Connector for HDFS. The higher this number, the better the performance will be. It is the most important factor for tuning. A value equal to the number of input data files is recommended. The number of data files creates an upper bound for the parallelism possible, so it is recommended to create multiple input data files in upstream applications if possible.

Default Directory (external table): Stores the log files and `.bad` files, which contain information about rows that could not be loaded. You can specify the same or different locations for **Log File** and **Bad File** creation.

Destination Connection: If you are using Oracle Wallet, specify credentials. Otherwise, the credentials from the SQL Developer connection will be used.

Delete Deployment Files on SSH Host: Controls whether the deployment files are deleted on the SSH host after the deployment is complete.

Summary

Test: Tests the operation and lets you see the DDL that will be generated.

Finish: Performs the operations that you specified.

6.151 Install/Run Oracle REST Data Services

This full installation wizard is displayed if you click Tools, then REST Data Services, then Install. The more limited run wizard (for running in standalone mode) is displayed if you click Tools, then REST Data Services, then Run.

The wizard is a convenient alternative to the command-line installation of Oracle REST Data Services that is documented in *Oracle REST Data Services Installation, Configuration, and Development Guide*. (However, to install and use Oracle REST Data Services, you should be familiar with the main concepts and techniques described in that guide.)

6.151.1 ORDS File Locations

ORDS WAR File Location: For the Oracle REST Data Services `ords.war` file, specify whether to use the one that is included with SQL Developer or one that you downloaded from the Oracle Technology Network.

Configuration Files Location: Specify a location to be used for Oracle REST Data Services configuration files. (This can be a location of your choice, and should be one dedicated to storing the configuration files.)

Optionally, reset the configuration files location using the value from the `ords.war` file that you specified.

6.151.2 Database Connection

Specify the database connection information for the database on which to install Oracle REST Data Services: host, port, and the database name (SID) or service name.

ORDS Public User: Specify the user name and password for the Oracle REST Data Services public user (the database user that has privileges to access the Oracle REST Data Services RESTful services). If you do not skip the verification of this user, you will be prompted for that user's name and password in the **Verify Oracle REST Data Services Schema** dialog box.

Skip verify/install of Oracle REST Data Service Schema: Causes the information that you entered on this page not to be verified. This option allows you to continue with the wizard steps if you do not know all the information.

6.151.3 Database Options

This step of the wizard appears only when installing Oracle REST Data Services into a multitenant container database (CDB).

Select the CDB and its PDBs on which to install Oracle REST Data Services.

Install all read only PDBs to install/upgrade: Select this option to install Oracle REST Data Services on all read-only PDBs.

6.151.4 Select Tablespace

Select the default tablespace and temporary tablespace for the ORDS_METADATA and ORDS_PUBLIC_USER users.

6.151.5 PL/SQL Gateway

Specify credentials (user name and password) for the PL/SQL Gateway user, which is responsible for executing stored PL/SQL stored packages and procedures on the specified database connection.

Verify Database User: Displays a dialog box asking for the privileged user (for example, SYS AS SYSDBA) and that user's password to be used to verify the user credentials that you entered. If you enter correct information for a suitably privileged account, the credentials are checked and a message is displayed. (Note that if you do not enter valid credentials specified user or users, you can still continue with the wizard steps; but to be able to use Oracle REST Data Services, you will need to provide valid credentials.)

6.151.6 APEX RESTful Services

Use Application Express RESTful Services: Enable this option if you want to be able to use the Oracle Application Express RESTful services. If you enable this option, enter credentials for the proxy user that has privileges to access the Application Express RESTful services and the user that has privileges to create and edit Application Express RESTful services.

Verify Database User: Displays a dialog box asking the for the privileged user (for example, SYS AS SYSDBA) and that user's password to be used to verify the user credentials that you entered. If you enter correct information for a suitably privileged account, the credentials are checked and a message is displayed. (Note that if you do not enter valid credentials specified user or users, you can still continue with the wizard steps; but to be able to use Oracle REST Data Services, you will need to provide valid credentials.)

6.151.7 Run Standalone Mode

Oracle REST Data Services includes a web server that enables you to run in standalone mode. Standalone mode is designed for use only in development and test environments, and is not supported for use in production environments.

Run in Standalone Mode when installation completes: Causes Oracle REST Data Services to run in standalone mode when the installation completes.

- If you select (check) this option, SQL Developer creates the Oracle REST Data Services configuration files and performs the installation, and does not terminate. You can view the Oracle Rest Data Services Log by selecting **View > Log**. If you want to terminate the Oracle REST Data Services process, you can right-click on the Oracle REST Data Services Log and select **Terminate**, or you can click the red square icon.
- If you do not select (uncheck) this option, SQL Developer creates the Oracle REST Data Services configuration files and performs the installation, but automatically terminates after that. You can view the Oracle Rest Data Services Log by selecting **View > Log**.

HTTP Port: HTTP port to use.

Application Express Static Resources Location: If you are using Oracle Application Express, specify the location of the Application Express static resources, including images and CSS files.

6.151.8 ORDS Users

You can create Oracle REST Data Services users that are allowed to perform certain Oracle REST Data Services operations using Oracle SQL Developer:

- **ORDS Administrator:** An administrator responsible for administering the Oracle REST Data Services configuration using SQL Developer.
- **ORDS RESTful Services User:** A provisioned user that allows SQL Developer to be used to create and edit Oracle REST Data Services RESTful services.
- **Requires Secure Sockets Layer:** Enable this option if your RESTful services require Secure Sockets Layer (SSL) security, for example, `https`.

6.151.9 Summary

Displays a summary of the information that you have specified. To change any information, press Back as needed and make changes, then press Next as needed to return to this page.

To perform the installation, click **Finish**. (Exactly what happens depends on your choice for the **Run in Standalone Mode when installation completes** option.) To cancel the installation, click **Cancel**.

After installation, the REST Data Services option appears in the Connections navigator display and in the context menu for the connection associated with the database on which Oracle REST Data Services is installed.

Note:

During installation, if you are already connected to the database on which you are installing Oracle REST Data Services, then you have to disconnect and reconnect to see the REST Data Services option.

6.152 Install SQL Translator

This dialog box is displayed if you right-click SQL Translators under SQL Translator Framework for a connection in the DBA navigator and select Install SQL Translator.

SQL Translator: SQL translator to be installed for the desired third-party SQL to be translated (for example, Sybase SQL Translator for Sybase).

Mode: online lets you perform the installation interactively in SQL Developer; **offline** creates a .zip file containing several files, including a Linux .sh file that you can run to perform the installation.

Create New Profile: If this option is checked, you can also create a new SQL translator profile associated with the translator being created. In this case, specify a name for the profile (for example, `sybase_profile`) and an existing Oracle database schema to own the profile. Whether you check this option or not, you can also create new profiles separately.

Click **Apply** to start the installation of the SQL translator. (This installation involves several steps, and it may take a long time.)

Related Topics:

[SQL Translators](#) (page 1-64)

[SQL Translator FrameWork](#) (page 1-64)

[Using DBA Features in SQL Developer](#) (page 1-54)

6.152.1 Installing a Translator and Creating a Profile: Usage Notes

Installing a SQL translator in an Oracle schema enables you to run third-party statements and procedures, have SQL Developer automatically translate them to Oracle syntax, and see the results you would expect from Oracle.

Install the SQL translator in an Oracle schema where the user has been granted the usual privileges for connecting and creating database objects, as well as the CREATE SQL TRANSACTION PROFILE privilege. For example:

```
GRANT CREATE SQL TRANSACTION PROFILE to USER3;
```

After you click **Apply**, SQL Developer checks if the database user has the necessary privileges; and if it does not, it prompts you to log in as the SYS user and then confirm the granting of certain privileges.

The actual translator installation can take a long time, during which the Online SQL Translator Install progress box is displayed. If you want to continue working in SQL Developer during the installation, you can click **Run in Background**.

After the installation, grant execute access on the translator to the database user. For example:

```
GRANT EXECUTE on SYBASE_TSQL_TRANSLATOR to USER3;
```

You can now right-click a profile that is associated with the translator and select Open SQL Worksheet with Profile. In the SQL worksheet that is displayed, you can run SQL statements and procedures in that third-party SQL. For example, in the worksheet for a profile associated with the SYBASE_TSQL_TRANSLATOR translator, you could enter the following Sybase syntax that is not recognized by native Oracle Database:

```
select top 5 * from my_table;
```

If you double-click the profile name in the DBA navigator and then select the SQL Translations tab, you can see the SQL text and the translated text for each statement executed. If you want, and if you are sufficiently knowledgeable, you can modify the

TRANSLATED_TEXT cells for selected statements to change the way they are interpreted and executed by Oracle.

6.153 Invalid Database Version

The feature you are trying to use is not supported with Oracle Database Release 10 and earlier releases.

To use the feature, you must connect to an Oracle database with a version number of 11.1 or higher.

6.154 Load Extension

If preferences for a feature area are not displayed, click the **Load Extension** button to enable the relevant preferences.

Some SQL Developer features are implemented as Oracle-supplied extensions. Some of these are not loaded initially by default, but only after you use the feature in the SQL Developer interface or after you click Load Extension. For example, Migration preferences become available if you open the Migration Projects window (View > Migration Projects) or if you click the Load Extension button in the Migration preferences pane (if the button is visible).

6.155 Load Keyboard Scheme

This dialog box is displayed when you select Load Keyboard Scheme from More Actions when specifying shortcut key preferences for SQL Developer. You can load a set of predefined key mappings for certain systems and external editing applications. If you load any preset key mappings that conflict with changes that you have made, your changes are overwritten.

You can specify Default to restore the shortcut key mappings to the SQL Developer defaults for your system.

6.156 LOB Descriptors

Use this dialog box to specify storage options for LOB (large object) columns, enabling you to override the default storage options. You can add a LOB descriptor or delete a selected LOB descriptor.

Columns: Use the shuttle buttons to move the column(s) you want from Available Columns to Selected Columns.

LOB Segment Name: Name for the LOB data segment.

Storage Enabled in Row: Select to enable storage in row. The LOB value is stored in the row (inline) if its length is less than approximately 4000 bytes minus system control information.

Chunk: Number of bytes to be allocated for LOB manipulation. If the value is not a multiple of the database block size, then the database rounds up in bytes to the next multiple. The maximum value is 32768 (32K), which is the largest Oracle Database block size allowed. The default CHUNK size is one Oracle Database block.

Pct Version: Maximum percentage of overall LOB storage space used for maintaining old versions of the LOB. The default value is 10, meaning that older versions of the LOB data are not overwritten until they consume 10% of the overall LOB storage space. You can specify a Pct Version value whether the database is running in manual mode (where it is the default) or automatic undo mode (where Retention is the

default). You cannot enter values for both a Pct Version value and the Retention option.

Freepools: Number of groups of free lists for the LOB segment, usually the number of instances in a Real Application Clusters environment or 1 for a single-instance database. You can specify this only if the database is running in automatic undo mode. You cannot enter both a Free Pools value and the Free Lists fields.

Retention: Check to keep old versions of this LOB column. You can only do this if the database is running in automatic undo mode and if you do not specify a Pct Version value.

Cache: Choose how the database should store blocks in the buffer cache.

Storage: Opens the Storage dialog box, where you can define storage options.

6.157 Maintain Spatial Metadata

This dialog box is displayed if you select a table and specify Update Spatial Metadata. Internally, SQL Developer uses the information in this box to modify the relevant row in the `xxx_SDO_GEOM_METADATA` and `xxx_SDO_INDEX_INFO` views..

Table Name: Name of a feature table that has a column of type `SDO_GEOMETRY`.

Spatial Index Name: Name of the spatial index associated with the specified combination of table and column names.

Spatial Column: Name of the column of type `SDO_GEOMETRY` in the specified feature table.

Coordinate System: The SRID (spatial reference identifier) value for the coordinate system for all geometries in the column, or null if no specific coordinate system should be associated with the geometries. For example, 8307 is the SRID value for the "WGS 84 longitude/latitude" coordinate system.

Spatial Dimensions: The dimensional information (DIMINFO column in `USER_SDO_GEOM_METADATA`). Note that for *geodetic* layers, you must specify the dimensional extents in the index metadata as -180,180 for longitude and -90,90 for latitude, and the tolerance value is interpreted as meters.

Name: Name of the dimension, for example, X or Longitude.

Lower Boundary: Lower-left point of the dimension. For geodetic data, must be -180 for longitude and -90 for latitude.

Upper Boundary: Upper-right point of the dimension. For geodetic data, must be 180 for longitude and 90 for latitude.

Tolerance: A level of precision with spatial data, reflecting the distance that two points can be apart and still be considered the same (for example, to accommodate rounding errors). For geodetic data, tolerance is specified as meters.

Related Topics

- [Spatial Support in SQL Developer](#) (page 1-71)

6.158 Manage Columns

Use this dialog box to reorder, hide, or show columns in the display when you are editing a table or data in a table. To move a column higher or lower in the display order, click the column and then click the appropriate icon on the left. To hide a

column or to show a hidden column, click the column if necessary and click the appropriate icon between the Shown Columns and Hidden Columns lists.

When you are finished, click **OK**. (Your action affects only the current display; it does not change the table definition or any data in the table.)

6.159 Manage Connections (REST)

This dialog box is displayed when you click Tools, then REST Data Services, then Manage Connections. It contains two tabs:

- **Administration:** Displays the Edit REST Data Services Connection dialog box, where you can add, edit, and delete database connections for Oracle REST Data Services administration.
- **Development:** Displays the RESTful Services Connection dialog box, where you can add, edit, and delete database connections for RESTful Services development. Connections that you create are available for selection when you click the connect icon in the REST Development navigator.

Connection Name: Name for the connection.

Username (Administration tab): The Oracle REST Data Services administrator that was created to be able to edit the REST Data Services administration.

Username (Development tab): If you are using Oracle REST Data Services RESTful Services, specify the Oracle REST Data Services user that you created that allows SQL Developer to access Oracle REST Data Services RESTful Services.

If you are using Application Express RESTful Services, specify the Application Express user that you created using Application Express.

http or https: Protocol to use for connecting (HTTP or HTTPS).

Hostname: The host where Oracle REST Data Services is running.

Port: The port number on which Oracle REST Data Services is listening. Example: 8080

Server Path: The context root where Oracle REST Data Services is deployed.
Example: /ords

Schema/Workspace (Development tab only): Schema or workspace to use for RESTful Services development.

- If you are using Oracle REST Data Service RESTful Services, specify the schema alias.
- If you are using Application Express RESTful Services, specify the Application Express workspace that the user is assigned to.

6.160 Manage Features and Updates

This dialog box is displayed when you click Tools, then Features. It lets you enable or disable features in the product.

Features tab

The Extensions tab determines which features SQL Developer uses when it starts. (Technically, the features are internally implemented as "extensions".) Some features are mandatory, and users cannot remove or disable them; some features are optional,

and you can enable or disable them.) If you change any settings, you must exit SQL Developer and restart it for the new settings to take effect.

To cause an extension not to be used at the next startup, uncheck its entry. (To completely remove an extension, you must go to `sqldeveloper\extensions` under the SQL Developer installation folder and delete the `.jar` files associated with that extension.)

For Version Control, the settings (selected or not, and configuration options if selected) affect whether the Team menu is displayed and what items are on that menu.

Check for Updates: Checks for any updates to the selected optional SQL Developer extensions, as well as any mandatory extensions. (If the system you are using is behind a firewall, see the SQL Developer user preferences for [Web Browser and Proxy](#) (page 1-121).)

Check for Updates: Automatically Check for Updates: If this option is checked, SQL Developer automatically checks for any updates to the selected optional SQL Developer extensions and any mandatory extensions at startup. (Enabling this option might increase SQL Developer startup time. You can manually check for updates by clicking **Help**, then **Check for Updates**.)

Search icon: Enter text to filter the Available Features field.

Clear Cache: Removes previously loaded features from the cache.

Installed Updates tab

Lets you see and remove installed updates. Click **Expand All** to expand all nodes and **Collapse All** to collapse all nodes in the tree.

6.161 Map Connection

Use this dialog box to change the Connection value for one or more objects in the currently selected cart.

Map Connection <connection-name> to Connection: Select the database connection to be used.

Apply Mapping to: Changes the Connection mapping for one of more objects in the currently selected cart.

- **Selected Object Only:** Changes the Connection value only for the selected object.
- **All Objects for Connection <connection-name>:** Changes the Connection value for all objects (and only those objects) with the specified connection name in the currently selected cart. If the cart contains objects from multiple connections, only the objects associated with the specified connection are affected.
- **All Objects for All Connections:** Changes the Connection value for all objects (regardless of their current Connection values) in the currently selected cart.

Related Topics

- [Using the Cart](#) (page 1-68)

6.162 Merge Migration Projects

This dialog box is displayed when you right-click a migration project and select Merge into Project.

The operation merges the catalog or catalogs of the selected project into the specified destination project.

- For databases with multiple catalogs (such as Microsoft SQL Server and Sybase), it merges the catalogs.
- For databases that do not use multiple catalogs (such as IBM DB2, Teradata, and MySQL), it moves schemas from the single source to the destination.

After the operation, a refresh of the source and destination projects may be required.

A possible use case is to perform a parallel load into the repository from multiple computers and/or user accounts. In this case, for performance use the same repository for multiple projects, and then later merge projects.

Related Topics

- [SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

6.163 Modify Pluggable State

Use this dialog box to modify the state of a pluggable database (PDB).

Database Name: Name of the PDB.

New State: Contains OPEN if the PDB is closed, or CLOSE if the PDB is open.

State Option: List of options relevant to the New State value: Read Write, Read Only, or Restricted for OPEN; Normal or Immediate for CLOSE.

- **Read Write:** Allows queries and user transactions to proceed and allows users to generate redo logs.
- **Read Only:** Allows queries but does not allow user changes.
- **Restricted:** The PDB is accessible only to users with the RESTRICTED SESSION privilege in the PDB.
- **Normal:** Is the PDB equivalent of the SQL*Plus SHUTDOWN command with the normal mode (waits for users to disconnect from the database).
- **Immediate:** Is the PDB equivalent of the SQL*Plus SHUTDOWN command with the immediate mode (does not wait for current calls to complete or users to disconnect from the database).

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.164 Modify Value

This dialog box is displayed when you right-click a variable in the Data or Smart Data pane during debugging and select Modify Value. You can modify the value for the selected data item (primitive value, string, or reference pointer) during debugging. Note: You cannot undo the action after you click OK, so be careful when making any changes.

Current Value: The value of the data item.

New Value: The new value for the data item (enter or select from a drop-down list).

- For a primitive value, you can enter a new value.
- For a reference pointer, you can enter the memory address of an existing object or array. To set a reference pointer to null, enter 0 as a memory address.
- For a string, you can enter either a new string value or the memory address of an existing string.

Interpret New Value as Object Address: If this option is checked, the New Value entry is interpreted as a memory address pointer to an object or array in the heap of the program you are debugging. For a string, this box must be checked check if the value you enter in the New Value field is the memory address of an existing string.

Related Topics

- [Running and Debugging Functions and Procedures](#) (page 1-36)
- [Debug a PL/SQL Procedure](#) (page 5-10) (tutorial)

6.165 Data Move Details

This dialog box specifies the source and target information for online data migration.

Source Connection: Database connection from which data is to be migrated.

Target Connection: Database connection to which data is to be migrated.

Converted Model: Converted model containing tables whose data is to be moved to the corresponding Oracle database tables.

Use qualified names from converted model for insert: If this option is checked, object names are qualified by the schema name.

Related Topics

- [SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

6.166 New/Edit Change Set

In a source control system, a change set lets you group files together in ways that help you group multiple files related to the same change. For example, if a bug fix involves edits to three files, you can create a new change set and add all these files to the change set. Then, as you work through the bug fix, the change set will help you track all related files for the same bug.

Name: Name for this change set.

Use Title as Comment: Causes the title of the change set to be used as the contents of the Comments box.

Comments: You can modify or replace any existing comments for this change set, for example, a reference to the bug tracking number for which the files in this change set apply.

Template or Previous Comments: Select from available comment templates, if your version control system uses comment templates.

6.167 New Procedure (Refactoring)

This dialog box is displayed if you are editing a procedure, select one or more PL/SQL statements, right-click, and select Refactoring, then Extract Procedure. The selected statements are encapsulated into the procedure to be created.

Defined Locally: For a standalone procedure, defines the newly refactored code in the definition section of the original procedure.

Stored: For a standalone procedure, defines the newly refactored code in a new standalone procedure.

Name: Name of the procedure to encapsulate the selected statements. For a packaged procedure, the newly extracted procedure text is placed immediately after the current procedure.

6.168 New/Edit Local or Remote Port Forward

When you create an SSH (Secure Shell), you can create one or more local port forwards and/or remote port forwards. You can also edit existing local or remote port forwards.

Name: A name to be given to the local or remote port forward.

Host: Host name for the forward.

Port: Port on the specified host.

Automatically assign [local | remote] port: Lets a port be assigned automatically.

Use specific [local | remote] port: Specify the port to be used.

Related Topics:

[Connections with SSH Authentication](#) (page 1-33)

6.169 Edit SSH Connection

This dialog box enables you to edit an SSH (Secure Shell) connection. For more information, see [Connections with SSH Authentication](#) (page 1-33).

Host: SSH server. SQL Developer will create an SSH session to this host, using the specified details.

Port: SSH port. The default port is 22.

Username: User name that will be used to authorize the SSH session.

Use Key File: Specifies that a key file should be used to provide authentication. The key file contains a private key that should correspond to a public key registered with the server. The server verifies that SQL Developer has access to the proper private key and thus the user is who he or she claims to be.

Key File: Path to the key file.

6.170 No Lock

This dialog box is displayed if multiple instances of SQL Developer are running on the system and if you attempt to open a unit test object, such as a test or suite, when that object is already open in another instance of SQL Developer. Thus, the current instance of SQL Developer is unable to get a lock on that object at this time.

Either use the other instance of SQL Developer to view or modify the object, or close the object in the other instance of SQL Developer and then open the object in this instance.

6.171 No Object Found

This dialog box is displayed if no objects could be found to satisfy the requested operation, such as trying to perform a "Describe" operation when the currently selected object is not valid for a SQL*Plus DESCRIBE statement.

6.172 No Object Selected

This dialog box is displayed if no object was selected for the requested operation, such as trying to perform a "Describe" operation when no object is selected in the SQL Worksheet.

6.173 Object Preferences

This dialog box lets you specify preferences for the display of data output during debugging. For a specified data type and its subclasses, you can control what to display in value columns, what to show when expanding the object (expressions, fields, or both), and (for fields) which fields to show and which to hide.

6.174 Offline Generation Error - Destination Directory

This dialog box is displayed if you specify a nonexistent directory or folder for the generation of offline data move files. Check the spelling of the path that you entered, or create the desired directory or folder, and try the operation again.

6.175 Open File

This is a standard box for selecting a file to open: use **Location** to navigate to (double-clicking) the folder with the file to open, then click the file to select it.

6.176 Oracle-Only Report

This dialog box is displayed if you select a non-Oracle (third-party) database connection for a report that applies only to Oracle database connections. Be sure to select an Oracle connection.

6.177 Paste

This dialog box is displayed if you click **Edit**, then **Extended Paste**. It shows a list of clipboard items, so that you can select the content to be pasted. Click **OK** to paste the selected content into the current location.

Clipboard Items: Clipboard items with content from copy operations. Usually displays the first line of the content.

Item Content: The content of the selected clipboard item.

6.178 Perforce: Add Files to Perforce

(Applies only if you have added support for Perforce.)

Use to begin the process of bringing a new file under Perforce control.

Files list: Lists the files to be added to Perforce.

Name: The names of the files.

Location: The physical location of the files in the Name column.

Use Changelist: Select a changelist to use when adding these files.

6.179 Perforce: Connection

(Applies only if you have added support for Perforce.)

Use to connect to the Perforce server.

Username: Enter the name by which you are known to Perforce.

Password: Enter your Perforce password, if Perforce has been set up to require passwords.

Port: Enter the port number set up for the Perforce server. If the Perforce server is on another machine, prefix the port number with the machine name and a colon.

Client: Click on the Browse button, then navigate to select your preferred client.

Connection name: The preferred Perforce connection (if you have more than one) for this session.

Test Connection: Click if you want to test that a connection can be made to the Perforce server. The results are displayed in the text area below.

6.180 Perforce: Create Changelist

(Applies only if you have added support for Perforce.)

Use to create a changelist, to submit changed files to the Perforce depot.

Connection: The preferred Perforce connection (if you have more than one) for this session.

Description: Describe the changelist you are creating. The description you create will be visible in the Pending Changelists window. From there you can submit the changelist to the Perforce depot.

Files: The name and path of the files to be selected for the changelist. Place a check in the box beside each file name.

Select All: Selects all files for addition to the changelist (places checks beside all files).

Deselect All: Selects no files for addition to the changelist (clears checks from all files).

You can also make annotations to files in your Perforce changelist. Perforce will store annotations as a comment linked to every file in the revision. When you modify these files later in Perforce, you can view the sequence of revisions or changelists to these files.

6.181 Perforce: Delete Files

(Applies only if you have added support for Perforce.)

Use to delete files immediately from the Perforce client workspace. Files that are deleted from Perforce are indicated by a black diagonal cross. If you want to retrieve files deleted from the client workspace, use the Perforce client. To do this, select Team > Perforce > Launch Perforce Client.

Files list: Lists the files that will be deleted.

Name: The names of the files.

Location: The physical location of the files in the Name column.

6.182 Perforce: Login

(Applies only if you have added support for Perforce.)

Use to log in to the Perforce server.

Server: Name of the computer that acts as your Perforce server. (Read-only field.)

Username: Your username on the Perforce server named in the previous field. (Read-only field.)

Password: Your password for the Perforce server.

If you do not have or cannot remember your username or password for the Perforce server, contact the member of your team who administers the server. The username and password for your Perforce server may or may not be the same as any other username/password combination you have. Make sure you are entering the correct combination.

6.183 Perforce: Open Files for Edit

(Applies only if you have added support for Perforce.)

Use to open a file or files to be edited, with the option of placing the file on the default changelist or a specified changelist.

Name: The names of the files to be edited.

Location: The drive and path of the files in the Name column.

Open on Changelist: Select the changelist from which you want to open the file. Automatically opens on the default changelist.

Sync files to head position before opening: Select if you want to bring the selected files up to date with the controlled latest version in the Perforce depot.

Lock Files: Select if you want to prevent other users from working on the file until you have finished with it.

6.184 Perforce: Revert Files

(Applies only if you have added support for Perforce.)

Use to revert files to their immediately preceding state.

Files list: Lists the files that will be reverted.

Name: The names of the files.

Location: The physical location of the files in the Name column.

6.185 Perforce: Sync Files

(Applies only if you have added support for Perforce.)

Use to synchronize files in your Perforce client workspace with the controlled versions in the Perforce depot.

Connection: The preferred Perforce connection (if you have more than one) for this session.

Files list: Lists the files that will be synchronized.

Name: The names of the files.

Location: The physical location of the files in the Name column.

Sync: Select the revision against which the files in the files list will be synchronized.

Head Revision: Select to synchronize to the latest controlled version in the Perforce depot.

Revision Number: Select to synchronize to a particular revision number in the Perforce depot, then enter the revision number in the accompanying Target box.

Changelist: Select to synchronize to a particular Perforce changelist, then enter the changelist name in the accompanying Target box.

Label Name: Select to synchronize to a particular label name, then enter the name of the label in the accompanying Target box.

Date: Select to synchronize to a particular date or date-and-time combination, then enter the date or date-and-time combination in the accompanying Target box. The format is the same as used by the Perforce client application, which is displayed in the dialog.

Force sync: Select to reinstate files from the Perforce depot, even when this will overwrite or remove files in your Perforce client workspace.

Preview only: Select to report the changes that would be made during synchronization, without actually synchronizing.

6.186 Perforce: Submit Changelist

(Applies only if you have added support for Perforce.)

Use to submit changelists containing edited files to the Perforce depot.

Description: Enter any descriptive comments to identify the changes you have made to the files you are submitting.

Files to submit: Lists the files that will be submitted to Perforce. Select and deselect individual files to include or exclude them.

Comments: Use to add comments to accompany the submit action. You will later be able to see these comments when viewing the history of the files.

Template or Previous Comments: Select from available templates, or from previous comments you have added to other files.

Select All: Selects all the files currently displayed in the Files list.

Deselect All: Deselects all files in the list.

6.187 Plug In Pluggable Database

Use this dialog box to plug in a PDB. Plugging in consists of creating a PDB based on specifications in the XML file from when the original PDB was unplugged.

Database Name: Name of the PDB.

AS CLONE: Specify only if the target CDB already contains a PDB that was created using the same set of data files. The source files remain as an unplugged PDB and can

be used again. Specifying AS CLONE also ensures that Oracle Database generates new identifiers, such as DBID and GUID, for the new PDB. If you specify AS CLONE, then you cannot specify the MOVE or NOCOPY clauses.

XML File Name: Specify the XML file containing the metadata for the PDB to be plugged in.

Source File Name Conversions: Determines how the database generates the names of files (such as data files and wallet files) for the PDB.

- **None:** The database first attempts to use Oracle Managed Files to generate file names. If you are not using Oracle Managed Files, then the database uses the PDB_FILE_NAME_CONVERT initialization parameter to generate file names.
- **Custom Names:** Select a Source Files/Target Files pair.
- **Custom Expressions:** Specify one or more Source File Expression/Target File Expression pairs. Each pair item is a string found in names of files associated with the seed (when creating a PDB by using the seed), associated with the source PDB (when cloning a PDB), or listed in the XML file (when plugging a PDB into a CDB).

Copy Action: Determines how the database generates the names of files (such as data files and wallet files) for the PDB.

- **Nocopy:** The files for the PDB remain in their current locations. Use this option if there is no need to copy or move the files required to plug in the PDB.
- **Copy:** The files listed in the XML file are copied to the new location and used for the new PDB. You can also specify Source File Names Conversions to use pattern replacement in the new file names.
- **Move:** The files listed in the XML file are moved, rather than copied, to the new location and used for the new PDB. You can also specify Source File Names Conversions to use pattern replacement in the new file names.

Storage: You can specify storage limits for the PDB total size or temporary tablespace usage, or both; or unlimited storage for either or both.

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.188 Print Preview

This dialog box is displayed if you click File, then Print Preview. It displays an image of the page or pages to be printed. You can click Page Setup and Print Options to specify options for printing, or click Print to select the printer and perform the printing.

Page Setup: Displays a dialog box where you can specify the following for any diagram print operations: media **Size** (Letter, Legal, or other predefined size) and **Source** (Automatically Select or a specified paper source), **Orientation** (Portrait, Landscape, Reverse Portrait, Reverse Landscape), and **Margins** (left, right, top, bottom).

Print Options: Displays a dialog box where you can specify options for the border, header and footer, text, and zoom.

6.189 Privilege Warning for Migration

This dialog box is displayed if you click Verify in the Quick Migrate box and the database user for the connection does not have all privileges necessary for a multischema migration. For multischema migrations, this user must be granted the RESOURCE role with the ADMIN option; and this user must also be granted the CREATE ROLE, CREATE USER, and ALTER ANY TRIGGER privileges, all with the ADMIN option.

If you are performing a single-schema migration, you can ignore this warning.

6.190 Publish to REST

Use the RESTful Services wizard to publish a query to REST and create a RESTful Service from the display grid. To do this, right-click the grid and select **Publish to REST**.

You need to install Oracle REST Data Services release 3.0.5 or later.

- **Module Name:** Name of the RESTful service module that is unique. Select an existing module or create a new one.
- **URI Prefix:** Base of the URI that is used to access this RESTful service. If you select an existing module, then the URI Prefix is automatically displayed.
- **URI Pattern:** A pattern for the resource template.
- **Pagination Size:** Default pagination for a resource handler HTTP operation GET method, that is, the number of rows to return on each page of a JSON format result set based on a database query. The default is 25 items for a page.
- **SQL:** Displays the SQL query used to obtain the grid.

Related Topics:

[RESTful Services Wizard and Object Dialog Boxes](#) (page 6-127)

6.191 Recent Files

This dialog box displays files recently opened in SQL Developer.

Files: A list of files opened in SQL Developer, with the most recent file first. The Show All option determines whether the list includes only files opened implicitly or files opened implicitly or explicitly.

Show All: If this option is checked, the list includes both explicitly and implicitly opened files; if this option is not checked, the list includes only implicitly opened files. Explicitly opened files are those that you opened directly; implicitly opened files are those that SQL Developer opened to support your work (for example, while you were debugging).

6.192 Delete or Truncate Repository

The Delete Repository dialog box is displayed if you click Tools, then Migration, then Repository Management, then Delete Repository or Truncate Repository.

Deleting a repository removes all schema objects that are used for the migration repository. Truncating a repository deletes all data from schema objects that are used

for the migration repository, but does not delete the schema objects themselves, effectively leaving you with an empty repository.

Repository: Name of the database connection in which to delete or truncate the migration repository.

6.192.1 Error Truncating Repository

If an error occurs during an attempt to truncate the migration repository, do either of the following:

- Delete the repository and create it again in the same schema.
- Change the schema: create a new schema for the repository, and delete the old repository.

6.193 Rename Local Variable

This dialog box is displayed if you right-click a variable name in the display of the source code for a function or procedure, and select Refactoring and then Rename Local Variable. Specify the desired new name for the variable.

6.194 Rename Tab

This dialog box is displayed if you right-click a tab name and select Rename. You can change the name of the tab to something more descriptive in this case, for example, changing *Script Output* to *Employees Query*.

The new name that you specify is used only for as long as the tab is displayed. If you close the tab and cause it to be displayed again, the default tab name appears.

6.195 Rename Procedure

This dialog box is displayed if you try to rename a procedure. Specify a unique new name for the procedure.

6.196 Repair (Validate) Oracle REST Data Services

The Oracle REST Data Services Repair Wizard enables you to check if the Oracle REST Data Services installation is valid. This is the same as running the `validate` command in the command-line interface. For more information, see the “Validating the Oracle REST Data Services Installation” section in *Oracle REST Data Services Installation, Configuration, and Development Guide*.

If Oracle REST Data Services is installed on a database schema before installing Application Express, then Oracle REST Data Services attempts to find the Application Express schema and create a view that joins the relevant tables in the schema to those in the REST schema (`ORDS_METADATA`). If the Application Express schema is not found, then Oracle REST Data Services will create a stub view for the missing tables. You can use the Repair Wizard to force Oracle REST Data Services to reconstruct the queries against the Application Express schema.

The Repair Oracle REST Data Services wizard is displayed if you click **Tools**, then **REST Data Services**, then **Repair**.

ORDS File Locations: For the Oracle REST Data Services `ords.war` file, specify whether to use the one that is included with SQL Developer or the one that you

downloaded from the Oracle Technology Network for repairing the Oracle REST Data Services schema.

Database Connection: Specify the database connection information for the database on which to install Oracle REST Data Services: host, port, and the database name (SID) or service name. Additionally, specify the user name and password for a user with SYSDBA privileges to uninstall Oracle REST Data Services (enter SYS credentials).

Database Options: This step of the wizard appears only for a multitenant container database (CDB). Select the CDB and its PDBs to validate. You can validate only those that have Oracle REST Data Services installed.

Repair Summary: Click **Finish** to start the validation process.

6.197 REST Data Services Connection

Use this dialog box to specify properties for an Oracle REST Data Services connection.

Connection Name: Name for the Oracle REST Data Services connection.

Username: Oracle REST Data Services user name. (For example, might be your Application Express user name.)

http or https: Protocol to use: `http` or `https` (secure socket layer)

Hostname: Host where the REST Data Services is running.

Port: Port number on which the Oracle REST Data Services is listening.

Server Path: Context root where the REST Data Services is deployed. Example: `/ords`

Related Topics

- [Oracle REST Data Services Administration](#) (page 3-1)

6.198 REST Data Services Connection: Authentication Failed

This topic explains what to do if you try to connect to Oracle REST Data Services as described in [Connecting to Oracle REST Data Services](#) (page 3-3), but encounter the Authentication Failed error. The error box content starts with “Cannot connect to ORDS” and includes `ValidatorException: PKIX path validation failed`.

In this case, you must download and install the necessary certificate. The procedure depends on whether you are using SQL Developer on a Windows system or on a Macintosh or Linux system.

On Windows Systems

1. Close SQL Developer.
2. Use a web browser to open any link to Oracle REST Data Services. For example, a URL in this format:

```
https://yourserver:port/ords/oauth/
```

3. Click **Certificate Error** and view the certificate.
4. In the Certificates dialog box, click the **Details** tab, click **Copy to File**, then click **OK**.
5. In the Certificate Export Wizard, on the Export File Format page, select DER encoded binary X.509 (.CER).

6. On the File to Export page, browse to choose a location and specify a name for the file. (The location and file name can be any that you want.) For example: D:\mycert.cer
7. Go to the bin folder under where your JDK is installed.
If you are not using SQL Developer with a bundled JDK, this is %JAVA_HOME%\bin.
If you are using SQL Developer with a bundled JDK, this is jdk\jre\bin under where you installed SQL Developer.
8. In Windows File Explorer, right-click the bin folder and select **CMD Prompt Here as Administrator**.
9. Recommended: Back up your KeyStore file before performing the next step.
10. Enter a command in the following format:

```
keytool -import -alias ENTRY_NAME -keystore PATH_TO_KEYSTORE -file PATH_TO_CERT
```

where:

- *ENTRY_NAME* is a name to alias the certificate in the store.
- *PATH_TO_KEYSTORE* is a location inside the SQL Developer or JDK installation.
- *PATH_TO_CERT* is where you exported the certificate file.

For example:

```
keytool -import -alias sqld1 -keystore D:\sqldeveloper\jdk\jre\lib\security\cacerts -file D:\mycert.cer
```

11. Start SQL Developer and try to connect again.

On Macintosh or Linux Systems

1. Close SQL Developer.
2. Use a web browser to open any link to Oracle REST Data Services. For example:
`https://yourserver:8443/ords/oauth/`
3. Drag the certificate image from the dialog box onto your desktop.
4. Determine your JAVA_HOME by entering the following command:

```
/usr/libexec/java_home
```

The result should be similar to the following:

```
/Library/Java/JavaVirtualMachines/jdk1.8.0_45.jdk/Contents/Home
```

5. Go to the location that is JAVA_HOME (from the preceding step) plus `/jre/lib/security`. For example

```
cd /Library/Java/JavaVirtualMachines/jdk1.8.0_45.jdk/Contents/Home/jre/lib/security
```

or (if \$JAVA_HOME is set)

```
cd $JAVA_HOME/jre/lib/security
```

6. Back up your KeyStore file. For example:

```
sudo cp cacert cacertbakp
```

7. Import the certificate. When prompted, enter the keystore password (which by default is `changeit`). For example:

```
sudo keytool -import -alias sqld1 -keystore cacerts -file /Users/youruser/
Desktop/yourserver.cer
Password:
  (Enter keystore password.)
Owner: CN=yourserver
Issuer: CN=yourserver
Serial number: 235b0f74a736100b
Valid from: Thu Apr 21 09:55:22 CDT 2016 until: Fri Apr 21 09:55:22 CDT 2017
Certificate fingerprints:
  MD5: AA:DE:74:49:A8:C5:85:BB:A8:74:D0:24:FF:00:CC:BC
  SHA1: E3:AA:B1:7C:A4:BB:C0:04:C5:AC:26:DD:DE:2F:95:33:66:28:EF:F6
  SHA256: 85:60:8A:AA:96:85:A8:9A:92:BE:57:CC:F0:E2:56:39:86:EF:CE:6E:9E:
74:FF:F7:C1:90:6B:5D:12:14:3B:53
  Signature algorithm name: SHA256withRSA
  Version: 3
Trust this certificate? [no]: yes
Certificate was added to keystore
```

8. Start SQL Developer and try to connect again.

6.199 REST Data Services Connection: Remove Connection

Use this dialog box to remove (delete) an Oracle REST Data Services connection.

Connection: Name of the Oracle REST Data Services connection to be removed.

Related Topics

- [Oracle REST Data Services Administration](#) (page 3-1)

6.200 RESTful Services Wizard (Auto-Enable REST Access)

This RESTful Services wizard lets you conveniently enable or disable Oracle REST Data Services access for the schema associated with a database connection, or for individual objects in a schema for which REST access is enabled -- if Oracle REST Data Services has been installed on the system containing the database with that schema. The wizard is displayed when you right-click the connection name or the name of a table or view in the Connections navigator, then select Enable REST Services.

For an explanation of using the AutoREST feature versus creating resource modules, as well as more information about this wizard, see [Automatically Enabling REST Access to a Schema, Table, or View \(AutoREST\)](#) (page 3-11).

Wizard pages:

- [Specify Details](#) (page 6-126)
- [RESTful Summary](#) (page 6-127)

6.200.1 Specify Details

Specify details for the operation.

Enable schema (if the schema was selected) or **Enable object** (if a specific table or view was selected): Controls whether Oracle REST Data Services access is enabled for the specified schema or object.

Schema alias (if the schema was selected) or **Object alias** (if a specific table or view was selected): Alias for identifying the schema or object.

Authorization required: For a schema, controls whether Oracle REST Data Services should require user authorization before allowing access to the Oracle REST Data Services metadata catalog of this schema. For an object, controls whether Oracle REST Data Services should require user authorization before allowing access to the Oracle REST Data Services metadata for this object. If this option is selected, the name of the privilege role required for a user to perform the access is displayed.

6.200.2 RESTful Summary

Displays a summary of the options that you have specified. To make any changes, click **Back** as needed to return to the desired page or pages.

To confirm and perform the requested operation, click **Finish**.

Related Topics

- [RESTful Services Wizard and Object Dialog Boxes](#) (page 6-127)
- [REST Development](#) (page 3-12)

6.201 RESTful Services Wizard and Object Dialog Boxes

The RESTful Services wizard lets you create a RESTful service; individual dialog boxes let you edit the properties of objects associated with a RESTful service.

Create REST Module, Template, and Handler using the Wizard

The RESTful Services wizard lets you create the resource module, resource template, and resource handlers for a RESTful service. To display the wizard, in the REST Development pane or in Connections navigator, expand the REST Data Services node, right-click Modules, and select **New Module**.

Note:

Only in REST Development, you can create the resource handler using the wizard. In Connections navigator, you need to right-click the template name and select **Add Handler** to create the resource handler.

Create REST Template using the Wizard

The RESTful Services wizard lets you create a resource template for a module. To display the wizard, expand REST Data Services node, expand Modules, right-click the module name, and select **Add Template**.

Create REST Handler

To create the REST Handler, expand the REST Data Services node, expand Modules, and select the Template name. Right-click the Template name, and select **Add Handler**. Select the required method (GET, POST, PUT, DELETE).

Edit REST Module, Template, or Handler

Right-click the REST module name, template name, or handler, and select **Edit**.

Open Browser

The Open Browser feature is available only when you use REST Data Services in the Connections navigator, which is applicable only for Oracle REST Data Services 3.0.5 and later.

When you click Modules, Privileges, or Roles in the Connections navigator, the browser is automatically displayed containing a list of objects (for example, modules, privileges, and roles) and their related information.

Clicking the module name, template name, or privilege name will automatically display the browser and the detail information.

Delete REST Module, Template, Handler, Privilege or Role

To delete a REST module, template, handler, privilege, or role, right-click its name and select **Delete**.

The wizard consists of the following steps:

- [Resource Module](#) (page 6-128)
- [Resource Template](#) (page 6-128)
- [Resource Handler](#) (page 6-129)
- [RESTful Summary](#) (page 6-130)

6.201.1 Resource Module

Specify properties of the resource module.

Module Name: Name of the RESTful service module. Case sensitive. Must be unique.

URI Prefix: Base of the URI that is used to access this RESTful service. Example: `hr /` means that all URIs starting with `hr /` will be serviced by this resource module. (The displayed example changes to include what you enter.)

Publish: Makes the RESTful service publicly available for use.

Pagination Size: Default pagination for a resource handler HTTP operation GET method, that is, the number of rows to return on each page of a JSON format result set based on a database query.

Origins Allowed: Origins that are allowed to access the resource templates. (Click the plus (+) to add each origin.) For example:

```
http://example1.org  
https://*.example2.com
```

Note:

After creating a resource module through the REST Development pane, you can **upload** it, which updates the RESTFUL Services definition on the server.

6.201.2 Resource Template

Specify properties of the resource template.

URI Pattern: A pattern for the resource template. For example, a pattern of `/objects/:object/:id?` will match `/objects/emp/101` (matches a request for the item in the emp resource with id of 101) and will also match `/objects/emp/`

(matches a request for the emp resource, because the :id parameter is annotated with the ? modifier, which indicates that the id parameter is optional). (The displayed example changes to include what you enter.)

A URI pattern can be either a route pattern or a URI template, although you are encouraged to use route patterns. Route patterns focus on decomposing the path portion of a URI into its component parts, while URI templates focus on forming concrete URIs from a template. For a detailed explanation of route patterns, see docs \javadoc\plugin-api\route-patterns.html, under <sqldeveloper-install>\ords and under the location (if any) where you manually installed Oracle REST Data Services.

Priority: Priority for the order of how the resource template should be evaluated (low through high).

HTTP Entity Tag (ETag): Identifies the type of entity tag to be used by the resource template. An entity tag is an HTTP Header that acts as a version identifier for a resource. Use entity tag headers to avoid retrieving previously retrieved resources and to perform optimistic locking when updating resources. Options include:

- **Secure HASH (default):** The contents of the returned resource representation are hashed using a secure digest function to provide a unique fingerprint for a given resource version.
- **Query:** Manually define a query that uniquely identifies a resource version. A manually defined query can often generate an entity tag more efficiently than hashing the entire resource representation.
- **None:** Do not generate an entity tag.

6.201.3 Resource Handler

Specify the properties of the resource handler. The specific options depend on the method type.

Method: HTTP request method for this handler: GET (retrieves a representation of a resource), POST (creates a new resource or adds a resource to a collection), PUT (updates an existing resource), or DELETE (deletes an existing resource). Only one handler for each HTTP method is permitted.

Requires Secure Access: Indicates whether the resource should be accessed through a secure channel such as HTTPS.

Source Type: Source implementation for the selected HTTP method:

- **Feed:** Executes a SQL query and transforms the results into a JSON Feed representation. Each item in the feed contains a summary of a resource and a hyperlink to a full representation of the resource. The first column in each row in the result set must be a unique identifier for the row and is used to form a hyperlink of the form: path/to/feed/{id}, with the value of the first column being used as the value for {id}. The other columns in the row are assumed to summarize the resource and are included in the feed. A separate resource template for the full representation of the resource should also be defined. Result Format: JSON
- **Media Resource:** Executes a SQL Query conforming to a specific format and turns the result set into a binary representation with an accompanying HTTP Content-Type header identifying the Internet media type of the representation. Result Format: Binary

- **PL/SQL:** Executes an anonymous PL/SQL block and transforms any OUT or IN/OUT parameters into a JSON representation. Available only when the HTTP method is DELETE, PUT, or POST. Result Format: JSON
- **Query:** Executes a SQL Query and transforms the result set into either a JavaScript Object Notation (JSON) or CSV representation, depending on the format selected. Available when the HTTP method is GET. Result Format: JSON or CSV
- **Query One Row:** Executes a SQL Query returning one row of data into a JSON representation. Available when the HTTP method is GET. Result Format: JSON

Data Format (if the Source Type is Query Source): JSON or CSV.

Pagination Size (GET handler): Default pagination size for the module's resource templates: size of the pagination window, or the number of rows to return for a database query.

MIME Types (PUT and POST handlers): Content types.

6.201.3.1 Resource Handler Editing Pane

When you edit a resource handler, the following tabs are available:

- **SQL Worksheet:** You can specify the query or PL/SQL code in the SQL Worksheet for the resource handler. The SQL Worksheet enables you to connect to the schema, execute the query or PL/SQL code, and view the results.

Note:

When the handler is added through the REST Data Services node in the Connections navigator, the Save REST Handler icon appears in the SQL Worksheet toolbar. Click this icon to save the RESTful Services definition directly to the database.

- **Parameters:** Any parameters for the query or the PL/SQL code.
- **Details:** You can edit the security and other attributes for the [Resource Handler](#) (page 6-129).

6.201.4 RESTful Summary

Displays a summary of the RESTful service properties that you have specified. To make any changes, click **Back** as needed to return to the desired page or pages.

To create the resource module, resource template, and resource handler as specified, click **Finish**.

Related Topics

- [RESTful Services Wizard and Object Dialog Boxes](#) (page 6-127)
- [REST Development](#) (page 3-12)
- [Installing Oracle REST Data Services](#) (page 3-1)

6.202 RESTful Services Connection

Use this dialog box to specify properties for a RESTful Services connection.

Connection Name: Name for the RESTful service connection.

Username: Application Express user name.

http or https: Protocol to use: `http` or `https` (secure socket layer)

Hostname: Host where the REST Data Services is running.

Port: Port number on which the Oracle REST Data Services is listening.

Server Path: Context root where the REST Data Services is deployed. Example: `/ords`

Schema or Workspace: If you are using Oracle REST Data Service RESTful services, specify the schema alias. If you are using Oracle Application Express RESTful services, specify the Application Express workspace to which the user is assigned.

Related Topics

- [REST Development](#) (page 3-12)

6.203 RESTful Services Privilege

Use this dialog box to specify properties for a RESTful Services privilege. You can create a RESTful Services privilege to restrict access to specified users for one or more RESTful Services modules. A RESTful Services privilege can include users belonging to one or more user groups.

To create a privilege, expand the REST Data Services node in the Connections navigator, right-click Privileges and select New Privilege. You can also create a privilege using the REST Development pane.

You can specify roles and protected modules to be associated with the privilege by using the arrow and double-arrow keys to move items between Roles and Selected Roles and between Modules and Protected Modules.

Name: Name for the RESTful service privilege.

Title: Name for the RESTful service privilege.

Description: Description for the RESTful service privilege.

Roles: Available roles that can be associated with this privilege.

Selected Roles: Roles to be associated with this privilege.

Modules: RESTful Services modules to which the privilege can be assigned.

Protected Modules: Modules to which this privilege is assigned.

Related Topics

- [REST Development](#) (page 3-12)

6.204 RESTful Services Roles

Use this dialog box to create and edit roles for RESTful Services. To create roles, expand the REST Data Services node in the Connections navigator display, right-click Roles and select New Role. To edit a role, right-click the role name and select Edit.

Role: Name for the RESTful Service role.

6.205 Select Current Repository

You can use this dialog box to reconnect to a migration repository after you have disconnected. In addition, if you have multiple migration repositories, and you can use this dialog box to switch from one to another.

Select Current Repository: Name of the database connection with the migration repository to be used for all operations relating to migrating third-party databases to Oracle.

6.206 Cannot Capture Table

This dialog box is displayed if you try to capture a third-party database before establishing and connecting to a current migration repository.

If no migration repository exists, create one by clicking **Tools**, then **Migration**, then **Repository Management**, then **Associate Repository**.

To make an existing migration repository the current one, right-click its connection in the Connections navigator and select **Associate Migration Repository**.

To open a connection to the migration repository, expand the node for its connection in the Connections navigator.

6.207 Report Panel

This panel is mainly used to display specific [SQL Developer Reports](#) (page 1-82). See the help for the relevant report or report type.

This panel is also used for displays of some other features, such as **Monitor Sessions** on the **Tools** menu.

6.208 Reset Expired Password (Enter New Password)

This dialog box is displayed if you attempt to create a new database connection or open an existing connection, and if the password associated with the used for the connection has expired. It is also displayed only if an OCI (thick) driver is available; if an OCI driver is not available, an error message is displayed instead of this dialog box.

To reset the password, enter the current password for the specified user, enter the new password, confirm the password (type the same new password), and click **OK**.

6.209 Revision Lister

This dialog box is displayed if you click List Revisions in the [Subversion: Branch/Tag](#) (page 6-140) dialog box. It contains a list of revisions in the repository.

Select the desired revision to use, and click **OK**.

6.210 Run/Debug/Profile PL/SQL

Use this box to specify parameter values for running, debugging, or profiling a PL/SQL function or procedure. (If you specify a package, select a function or procedure in the package.) A profile operation runs the function or procedure and collects execution statistics; it also requires auxiliary structures in the user schema.

Comment (Profile only): Descriptive comment to be included in the execution profile.

Target: Name of the function or procedure to run or to run in debug mode. (You have a choice only if you specified a package that has more than one subprogram.)

Parameters: List of each parameter for the specified target. The mode of each parameter can be IN (the value is passed in), OUT (the value is returned, or IN/OUT (the value is passed in, and the result of the function or procedure's action is stored in the parameter).

PL/SQL Block: A block of PL/SQL code created by SQL Developer. You should change the formal IN and IN/OUT parameter specifications in this block to actual values that you want to use for running or debugging the function or procedure.

For example, to specify 10 as the value for an input parameter named `in_rating`, change `IN_RATING => IN_RATING` to `IN_RATING => 10`.

When you click **OK**, SQL Developer runs the function or procedure.

If you are debugging a function or procedure, the debugging toolbar and one or more windows for debug-related information are displayed.

Related Topics:

[Using the PL/SQL Hierarchical Profiler](#) (page 1-40)

[Running and Debugging Functions and Procedures](#) (page 1-36)

[Debug a PL/SQL Procedure](#) (page 5-10)

6.211 Create/Edit Breakpoint

Use this box to create or edit a breakpoint to use when debugging a PL/SQL function or procedure.

6.211.1 Definition tab

Specify the definition of the breakpoint.

Breakpoint Type: Type of breakpoint, indicating when the breakpoint will occur. Options include breaking when one of the following occurs: a specific line of code (Source); exception class or other class; method, file, or watch.

Breakpoint Details: Options depend on the breakpoint type.

Breakpoint Group Name: Breakpoint group in which to include this breakpoint. Breakpoint groups can be edited, enabled, and disabled.

6.211.2 Conditions tab

Specify any conditions that apply to the breakpoint.

Condition: A SQL condition (WHERE clause without the WHERE keyword) restricting when the breakpoint occurs. For example, to specify that the condition should occur only when `status_code` is greater than 10, specify:

```
status_code > 10
```

Thread Options: You can specify whether the breakpoint occurs for all threads, or only when the breakpoint is hit by threads that either do or do not have a specified name.

Pass Count: The number of times the debugger should allow execution to pass over the breakpoint before the breakpoint occurs.

6.211.3 Actions tab

Specify the actions to be taken when the breakpoint occurs. The options you specify override any default values on the Debugger: Breakpoints: Default Actions pane for SQL Developer preferences.

Halt Execution: Pauses execution when the breakpoint occurs.

Beep: Beeps when the breakpoint occurs.

Log Breakpoint Occurrence: Sends a message to the log window when the breakpoint occurs. You can also specify the following to be included in each display: a tag, and a condition to be evaluated.

Enable/Disable a Group of Breakpoints: Enables or disables the specified breakpoint group when this breakpoint occurs.

Related Topics

- [Running and Debugging Functions and Procedures](#) (page 1-36)
- [Debug a PL/SQL Procedure](#) (page 5-10)

6.212 Save/Save As, or Select File

This is a standard box for saving information to a file or for selecting a file: use **Location** to navigate to (double-clicking) the folder in which to save or open the file, then specify the file name (including any extension) and, if necessary, the file type.

6.213 Save or Open Cart Tool Configuration

In dialog boxes for performing operations on the contents of a tab in the Cart window (such as export objects, "diff" objects, and copy objects), you can save an XML file that reflects the current settings in the dialog box, and you can later open that XML file to have its settings be the default values in the dialog box.

For example, if you have settings in the [Export Objects](#) (page 6-85) dialog box that are different from the SQL Developer default values, but you want to be able to use your customized settings as defaults for certain Cart export operations, you can save the customized settings in a Cart export configuration file, and later open that file for Cart export operations.

File (for Save Tool Configuration): Location and name of the XML file to contain the configuration settings. The default file name reflects the type of operation, such as `export_cart.xml`, `diff_cart.xml`, or `copy_cart.xml`.

Encoding (for Save Tool Configuration): Character set used for encoding of the XML data.

Open Cart (for Open Tool Configuration): Location and name of the XML file containing settings to be used as default values for this use of the dialog box.

6.214 Save Files

This box asks if you want to save the specified files before another action occurs (for example, saving procedures you had been editing before disconnecting).

6.215 Unable to Save Files

This box informs you that SQL Developer is unable to save the specified file or files. To cancel the attempt to save the files and to return to edit the relevant object, click **Cancel**.

6.216 Save Style Settings

This dialog box is displayed when you click **Save As** in the Code Editor: PL/SQL Syntax Colors pane when setting SQL Developer preferences. You can save the specified color settings as a named color scheme, which adds it to the drop-down list for **Scheme** in that pane.

Related Topics:

[Code Editor](#) (page 1-94)

[SQL Developer Preferences](#) (page 1-92)

6.217 Schema Differences Source or Destination Error

This error box is displayed if you click **Apply** before specifying the source or the destination, or both, for a schema differences operation.

Click **OK** to close the error box, then follow the instructions for performing the schema differences operation, as explained in [Database Differences](#) (page 6-75).

6.218 Script Execution Failed

This error box is displayed if the script generated by the Quick Migrate procedure fails before it completes its execution. The Build pane displays the error that caused the failure.

To close the error box and open the script in a SQL Worksheet window, where you can edit the text and run the corrected script, click **Yes**; or to close the error box without opening the script in a SQL Worksheet window, click **No**.

6.219 Script Generation Complete

This information box is displayed after you generate the controlling script and related files for performing an offline capture of a third-party database.

Click **OK** to close the box. Later, run the controlling script to generate output containing the converted model.

Related Topics:

[Offline Capture](#) (page 2-17)

6.220 Set Data Mapping

This dialog box is displayed if you right-click a captured model and select **Set Data Mapping**. You can use this dialog box to specify source data type mappings when migrating the specified third-party database to Oracle. If you are editing an existing mapping, you can change only the Oracle data type, precision, and scale information.

Show only data types used in the source model: If you check this option, only data types used in the selected captured model are shown. If you do not check this option, all valid data types for the source (third-party) database are shown.

Source Data Type: Data type name in the third-party database.

Oracle Data Type: Data type name in Oracle Database.

Type: *System* for a system-defined data type, or *User* for a user-defined data type.

Add New Rule: Displays the Add/Edit Rule dialog box, for specifying a mapping for another data type.

Edit Rule: Displays the Add/Edit Rule dialog box, for editing the selected mapping.

Remove Rule: Deletes the selected mapping.

Related Topics:

[Add/Edit Rule](#) (page 6-136)

[SQL Developer: Migrating Third-Party Databases](#) (page 2-1)

6.221 Add/Edit Rule

This dialog box is displayed if you click Add New Rule or Edit Rule in the [Set Data Mapping](#) (page 6-135) dialog box, which is used for specifying source data type mappings when migrating a specified third-party database to Oracle.

Source Data Type: Data type name in the third-party database.

Oracle Data Type: Data type name in Oracle Database.

Precision and Scale: Precision and scale values to be used for the source data type and Oracle data type during the conversion.

6.222 Set Encoding

This dialog box is displayed if you right-click a connection and select Set Encoding. Specify a character set for the connection. The character set that you choose is applied to the encoding of files under version control through that connection.

Platform Default (Newline Conversions): Uses the character set specified for the platform/operating system. Newline conversions for files crossing different platforms are handled automatically.

IDE Global Setting: Uses the default character set for the integrated development environment (IDE).

Other: Uses the selected character set.

Related Topics

- [Using Versioning](#) (page 1-53)

6.223 Set Pause Continue

This dialog box is displayed if you enter the SQL*Plus statement SET PAUSE ON in the SQL Worksheet and then run the worksheet contents as a script. After the SET PAUSE ON statement is processed, execution pauses (and this dialog box is displayed) after each statement until the SET PAUSE OFF statement is processed.

To have execution continue at the next statement, click **OK**.

6.224 Shared Wizard Pages

These pages are shared by two or more wizards, such as Database Diff, Database Copy, Database Export, Data Pump Export, and Data Pump Import. The page titles may vary slightly depending on the wizard.

- [Types](#) (page 6-137)
- [Specify Objects](#) (page 6-137)
- [Data](#) (page 6-137)
- [Summary](#) (page 6-138)

6.224.1 Types

(The page title is *Types to Diff*, *Types to Export*, or *Object Types*.)

Check the types of objects that you want to be included in the operation. You can click **Toggle All** to check and uncheck all individual types.

You must select at least one object type.

6.224.2 Specify Objects

You can limit the types or objects, and the objects within selected types, for the operation. You can specify a **Name** string filter (including % wildcard characters) and optionally an object type (or *All* for all object types), and click **Lookup** to display objects that satisfy the filter.

Name: Shows only objects that satisfy the filter specification. You can use the standard % wildcard character (for example, *EM%* to show all objects starting with *EM*). You can use quotation marks (double quotes) to specify mixed-case searching.

%: If this option is checked, the % wildcard is automatically added to any Name string that you specify.

More and Less: Click **More** to enable filtering by Schema and Type; click **Less** to hide the Schema and Type controls.

Schema: Select the database schema in which to search for objects.

Type: Select All Objects for all object types, or a specific type of object to further filter the display of objects.

Lookup: Click Lookup to display a list of objects that meet the selection criteria for the selected connection. Use the arrow keys to move selected objects or all objects from the available objects box to the selected objects box.

6.224.3 Data

(The page title is *Specify Data* or *Table Data*.)

You can limit the data for the operation by selecting objects in the same way as on the [Specify Objects](#) (page 6-137) page.

You can select the columns included in the operation by clicking the pencil icon in the Columns cell. (Note: For an export operation involving CLOB columns, only the first 32 KB of any CLOB is exported.)

You can limit the rows in the operation by entering a valid WHERE clause for a query on the table, without the WHERE keyword. For example, to restrict the data to rows where a column named RATING contains a value greater than 5, specify `rating > 5` and click **Go** to see the results of the query

6.224.4 Summary

(The page title is *Summary*, *Copy Summary*, *Diff Summary*, or *Export Summary*.)

You can review a hierarchical display of your specifications for the operation.

To go back and make any changes, click **Back**.

To perform the operation, click **Finish**.

6.225 Sign In (checking for updates)

This dialog box is displayed if any of the updates that you selected during the check for updates process are on a remote site that requires you to log in. Currently, all updates are on the Oracle Technology Network (OTN), so you must enter your OTN user name and password.

User Name: Your user name at the remote site.

Password: Your password at the remote site.

Sign Up: If you do not have an account at the remote site, click this link.

Find Password: If you have an account at the remote site but cannot remember your password, click this link.

6.226 Single Record View

The main use for this box, which is displayed by right-clicking the display grid for an object and selecting Single Record View, is to edit data for a table or view, one record at a time. After you change data in any cells in a row, you can apply the changes by clicking Apply or by navigating to another record. (For non-Data grids, the cells are read-only.)

Navigation icons: First (<<) moves to the first record, Previous (<) moves to the previous record, Next (>) moves to the next record, and Last (>) moves to the last record.

Apply: Applies changes made to the current data record.

Cancel: Cancels changes made to the current data record, and closes the box.

6.227 Save Snippet (User-Defined)

Use this box to create a user-defined snippet.

Category: Existing or new category in which to place the snippet. To create a new (user-defined) category, type the category name instead of selecting a category name from the list.

Name: Name of the snippet, as it will be displayed when users see the list of available snippets in the specified category. If an existing Oracle-supplied snippet has the same name in the same category, the user-defined snippet definition replaces the Oracle-supplied definition.

ToolTip: Optional tooltip text to be displayed when the mouse pointer stays briefly over the snippet name in the display of snippets in the specified category.

Snippet: Text that will be inserted for this snippet.

Related Topics

- [Using Snippets to Insert Code Fragments](#) (page 1-50)
- [User-Defined Snippets](#) (page 1-51)

6.228 Edit Snippets (User-Defined)

This box displays any existing user-defined snippets, and enables you to add, edit, or delete user-defined snippets.

To edit an existing user-defined snippet, select its row and click the **Edit User Snippet** icon, which displays the Save Snippet (User-Defined) dialog box.

To create a new user-defined snippet, click the **Add User Snippet** icon, which displays the Save Snippet (User-Defined) dialog box.

To delete a user-defined snippet, select its row and click the **Delete User Snippet** icon.

Related Topics

- [Using Snippets to Insert Code Fragments](#) (page 1-50)
- [User-Defined Snippets](#) (page 1-51)

6.229 Show SQL

This displays SQL statements reflecting the current context and content (if applicable).

6.230 Start Date, End Date, Repeat Interval, Comments (Scheduler)

Dialog boxes are available for specifying start and end dates, a repeat interval, and optional comments for an Oracle Scheduler job.

Start Date, End Date

For Start Date and End Date, specify the date and time, or click **Set Today** to specify the current date and time.

Tabs are provided for Frequency and various refinements of the frequency. Use as many as you need for specifying when to run the job.

Repeat Interval

Frequency tab: Specify a frequency for the job (such as DAILY or WEEKLY),

Other tabs (most in the form "By xxxx"): Select (check) **Enable** to be able to specify the value or values relevant to that frequency category.

A box at the bottom displays the keywords and values for the statement that SQL Developer will use to call the DBMS_SCHEDULER.CREATE_SCHEDULE procedure.

Comments

Optional descriptive comment about the scheduler object.

Related Topics

- [Scheduling Jobs Using SQL Developer](#) (page 1-67)

6.231 Subversion: Add Property

Use this dialog box to add a versioning property for the currently selected file or folder.

Related Topics

- [Using Versioning](#) (page 1-53)

6.232 Subversion: Add to Source Control

Use this dialog box to bring a new file under Subversion control.

Files List: Lists the names and physical locations of the files that will be added to Subversion.

Related Topics

- [Using Versioning](#) (page 1-53)

6.233 Subversion: Apply Patch

Use this dialog box to apply a previously generated patch. A patch must be applied to the same revision/tag from which it was generated.

The name and location are displayed for the project or set of files to which the patch will be applied.

Patch Source: Specify where the patch will be obtained from: the system clipboard or a file.

System Clipboard: Obtains the patch from the system clipboard.

File: Obtains the patch from a file. If you select this option, a suggested location and name for the patch file is already entered in the text box. You can change this by typing a new location and name into the box, or by using the **Browse** button to find a new location

Related Topics

- [Using Versioning](#) (page 1-53)

6.234 Subversion: Branch/Tag

This dialog box is displayed when you right-click a remote directory in the Subversion repository and select Branch/Tag. Create a branch by copying the current working copy or a revision from the repository to a selected location in the repository.

From: Location of the working copy or revision.

Working Copy: Causes the current working copy to be copied.

HEAD Revision: Causes the HEAD revision (the latest revision in the repository) to be copied.

Use Revision: Causes the revision specified in the text box to be copied. To see a list of revisions from which you can choose, click **List Revisions**.

To: Destination location.

Comment: Optional descriptive comment.

Switch to new branch/tag: If this option is checked, the existing working copy is switched to the new branch.

After you click **OK**, the SVN Console - Log pane is displayed at the bottom, with messages about commands that were executed.

Related Topics

- [Using Versioning](#) (page 1-53)

6.235 Subversion: Check Out from Subversion

Use this dialog box to check out modules from a Subversion repository, to create the initial local copies of files and folders that are being stored within a Subversion repository. It is these local copies, stored outside the Subversion repository, that you work on. This location and the files within it constitute the Subversion "working copy".

Note: With Subversion, there is no "check in" procedure, so you do not check in files that you have updated and check them out again when you next want to work on them. Instead, when you have finished working on your local copies of files, you **commit** them to the Subversion repository to bring the files held there up to date. Also, if you want to incorporate into your local copies changes that other developers have committed to the Subversion repository, you **update** them.

Destination: Directory or folder into which to place the checked out files. If this destination is not empty, a warning message will be displayed asking if you are sure you want to check out into this directory. (Attempting to check out files into a non-empty destination might reflect a mistake in specifying the destination, or it might be your intention.)

Use Revision: If this option is checked, the revision you specify in the text box is used. To see the available revisions, click the binoculars icon.

Depth: The level of recursion for selecting files to be checked out, from Infinity (all children at all levels under the selected item in the Versioning browser hierarchy) through Empty (only this item and no children).

Related Topics

- [Using Versioning](#) (page 1-53)

6.236 Subversion: Commit Resources

Use this dialog box to commit individual files to the Subversion repository. If a file is a child of a changed parent that has not been committed, you must either first commit the parent or instead commit the working copy.

The committed files will replace those in the repository as the most recent. Other developers who subsequently check out or update from these files will see the file changed in comparison with the previous version held in the repository.

Files List: Lists the names and physical locations of the files that will be committed to the Subversion repository.

Keep Locks: Retains the locks that you previously obtained on the files that you are about to commit. This will mean that other developers will still not be able to commit changes they may have made to the files.

Comments: Comments to accompany the commit action. You will later be able to see these comments when viewing the list of versions of a particular file.

Template or Previous Comments: A template with comment text for the Comments box.

Related Topics

- [Using Versioning](#) (page 1-53)

6.237 Subversion: Commit Working Copy

Use this dialog box to commit the working copy to the Subversion repository. The committed files will replace those in the repository as the most recent. Other developers who subsequently check out or update from these files will see the file changed in comparison with the previous version held in the repository.

Files List: Lists the names and physical locations of the working copy that will be committed to the Subversion repository.

Keep Locks: Retains the locks that you previously obtained on the files that you are about to commit. This will mean that other developers will still not be able to commit changes they may have made to the files.

Comments: Comments to accompany the commit action. You will later be able to see these comments when viewing the list of versions of a particular file.

Related Topics

- [Using Versioning](#) (page 1-53)

6.238 Subversion: Confirm Checkout

This dialog box is displayed if you attempt to check out from the repository root, as opposed to from the branches, tags, or trunk of the repository. To proceed with the checkout from the root, click **Yes**; to cancel this request, click **No**.

Skip This Message Next Time: If you enable this option, on future requests to check out from the repository root, this dialog box will not be displayed and the operation will proceed as if you had clicked Yes.

Related Topics

- [Using Versioning](#) (page 1-53)

6.239 Subversion: Create Remote Directory

Use this dialog box to create a remote directory for a connection in a Subversion repository.

Directory Name: Directory name to be associated with the specified URL.

Comments: Optional descriptive comment.

Related Topics

- [Using Versioning](#) (page 1-53)

6.240 Subversion: Create Subversion Repository

This information applies to creating a Subversion repository. A connection to the repository will be created automatically.

Repository Path: Location for the new Subversion repository. You can **Browse** to select the location.

File System Type: Data storage system type for the repository. For information about choosing a system, see "Version Control with Subversion" at <http://svnbook.red-bean.com/>.

- **Native:** The file system type being used by the operating system.
- **Berkeley DB:** Causes a Berkeley DB database to be used as the data storage system.

Connection Name: Name for this connection. If you leave this box blank, the connection will be given a name based on the URL of the repository location.

Related Topics

- [Using Versioning](#) (page 1-53)

6.241 Subversion: Create/Edit Subversion Connection

This information applies to creating or editing a Subversion connection. For information about SQL Developer support for versioning and Subversion, see [Using Versioning](#) (page 1-53).

Repository URL: Full, valid URL for the location of the Subversion repository. The following are URL schemas and the access methods they map to:

- `file:///` -- Direct repository access (on local disk)
- `http://` -- Access via WebDAV protocol to Subversion-aware Apache server
- `https://` -- Same as `http://`, but with SSL encryption
- `svn://` -- Access via custom protocol to an svnserve server
- `svn+ssh://` -- Same as `svn://`, but through an SSH tunnel

Connection Name: Name for this connection. If you leave this box blank, the connection will be given a name based on the URL of the repository location.

User Name: User name known to the repository, if the repository requires user and password validation.

Password: Password for the specified user, or blank if a password is not required.

Test Read Access: Attempts to establish a connection for read access to the Subversion repository.

Status: Displays the result of the test (success or an error message).

Related Topics

- [Using Versioning](#) (page 1-53)

6.242 Subversion: Delete Resources

Use this dialog box to delete the selected directory (folder) in the repository.

Comment: Comment explaining the deletion.

Related Topics

- [Using Versioning](#) (page 1-53)

6.243 Subversion: Edit Configuration File

This dialog box is displayed if you click Edit "server" in the [Versioning: Subversion](#) (page 1-117): General preferences pane. You can modify the Subversion configuration file directly.

Reset: Discards any changes that you have made and leaves the dialog box open.

To save any changes and close the box, click **OK**; to discard any changes and close the box, click **Cancel**.

Related Topics

- [Using Versioning](#) (page 1-53)

6.244 Subversion: Export Files

Use this dialog box to copy files from the Subversion repository to a local file system directory, or to copy working copies to a local file system directory.

Working Copy Path: The location of the files that will be copied for export. Only files that are under Subversion control will be exported.

Destination Path: A path that includes the directory where you want the files to be copied to.

Related Topics

- [Using Versioning](#) (page 1-53)

6.245 Subversion: Export Subversion Connections

Use this dialog box to export the details of one or more current Subversion connections to a file. The details can subsequently be imported from the file to re-create the connections.

File Name: The location and name for the file that will contain the connection details, or browse to a file/location using the **Browse** button.

Connections: Select one or more connections whose details will be exported.

Related Topics

- [Using Versioning](#) (page 1-53)

6.246 Subversion: Ignore

Use this dialog box to mark a file, or a pattern that identifies common file names, as content that Subversion should ignore. (This dialog box sets the `svn:ignore` property for the specified content.)

Often, a directory contains files that should not be kept under version control. For example, log files from a debug or batch operation do not need to be tracked or merged, yet they are often in the same directory as the shared code for a project. Such files should be marked to be ignored by Subversion.

Related Topics

- [Using Versioning](#) (page 1-53)

6.247 Subversion: Import Subversion Connections

Use this dialog box to import the details of Subversion connections from a previously created file.

File Name: The location and name for the file that contains the connection details, or browse to a file/location using the **Browse** button.

Connections: Select one or more connections whose details will be imported. If a connection to be imported already exists with the same URL, you will be asked to confirm whether you want to overwrite the existing connection details with the details in the imported connection.

Related Topics

- [Using Versioning](#) (page 1-53)

6.248 Subversion: Import to Subversion

Use this wizard to import source files into the Subversion repository To go from one step to the next, click Next; to go back to the previous step, click Back.

6.248.1 Destination

Use to identify the Subversion repository, and directory within the repository, where the imported files will be stored.

Repository Connection: The connection for the Subversion repository in which you want to store the imported files.

Path: The directory within the Subversion repository for storing the imported files.

6.248.2 Source

Source Directory: The directory containing the source files that you want to import into Subversion. Initially contains a path based on the item that was selected when you launched the wizard.

Comments: Comment text to accompany the imported files. The comments are recorded with the files in the Subversion repository and will be viewable with the version history of the files. You must enter some comment text; otherwise, an error will occur when you click Finish to attempt to perform the import operation.

6.248.3 Filters

Filters that will be applied to the import operation. If you do not want one or more of the filters to be applied, move them from Selected Filters to Available Filters using the left arrow keys. If necessary, you can use the right arrow keys to move filters from Available Filters to Selected Filters.

New: Displays a dialog box in which you can create a new filter that will be applied to the import operation. New filters are added to the Selected Filters list.

6.248.4 Options

You can configure options specific to the import operation.

Do Not Recurse: If this option is enabled, it prevents files being imported from directories subordinate to the one you identified on the Source page.

Perform Checkout : If this option is enabled, the imported source files will be checked out after import.

6.248.5 Summary

Displays the selected options for the import operation. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [Using Versioning](#) (page 1-53)

6.249 Subversion: Lock Resources

Use this dialog box to perform a Subversion lock operation on one or more checked out files (working copies).

Files List: Lists the names and physical locations of the files to be locked. You can individually select and deselect files.

Steal Lock: Breaks any existing locks and relocks the files for your use. Causes the `--force` option to be added to the underlying `svn lock` command.

Comments: Comments to accompany the action.

Related Topics

- [Using Versioning](#) (page 1-53)

6.250 Subversion: Merge

A merge operation copies changes made in one branch to another branch, or copies changes from a branch to the trunk (main line of development). It is typically used to bring another developer's work into your own files, and to merge private development back into the main line of development.

The merge is created by comparing the content of two revisions within the Subversion repository, and applying the differences to a Subversion working copy. If you subsequently want to use the result of the merge in the main line of development, you commit the working copy to the Subversion repository in the usual way.

Specify the following:

- **Merge Type**: *Merge Selected Revision Range*, *Reintegrate a Branch*, or *Merge Two Different Trees*.
- **Merge Resource**
- **Merge Options**

Your selection for Merge Type affects the content of subsequent displays, which can include the following.

From URL and its (start) revision to merge: The resource that is the basis of the comparison. (The resource entered in the To URL box will be compared against the resource entered here.)

HEAD Revision from Repository: Causes the comparison to be made against the most recently committed resources in the Subversion repository.

Use Revision: Causes the comparison to be made against a resource in the Subversion repository with a particular revision number. When selected, the accompanying text box becomes available. You can then enter a revision number into the text box, or click the List Revisions button to select the revision that you require.

To URL and its (end) revision to merge: The resource that will be compared with the base resource selected in the From URL box.

Same as "From" URL: Uses the same base repository location for both elements of the comparison.

Ignore Ancestry: Ignores any relationships between resources in the Subversion repository when comparing the start and end revisions. The effect of this will be to retain resources that have names identical to those they are being compared with, even though the resources have no common ancestry. The alternative is that a resource that predates an identically named one may be deleted and replaced with the later resource.

Dry Run: Causes the comparison to be performed without the changes being applied to the Subversion working copy. The results of the comparison are displayed in the Messages - Log window.

Related Topics

- [Using Versioning](#) (page 1-53)

6.251 Subversion: Pending Changes

This window shows files that have been added, modified or removed, either locally or remotely; files whose content conflicts with other versions of the same file; and files that have not been added to source control. This window is opened automatically when you first initiate an action that changes the local source control status of a file. You can also open this window manually.

The window shows any **outgoing changes** (files that have been added, modified or removed locally, and local files whose content conflicts with remote files), **candidates** (files that have been created locally but not yet added to source control), and **incoming changes** (files that have been added, modified or removed at a remote location).

You can restrict the files shown in the Pending Changes window by selecting a **scope** from the drop-down list. The default scope is Active Application, which will show files from the application currently selected in the navigator.

You can carry out appropriate source control actions on the files listed in the window, using the buttons in the toolbar. Clicking a button will either immediately initiate the operation, or open a dialog box through which you can choose options before proceeding.

Related Topics

- [Using Versioning](#) (page 1-53)

6.252 Subversion: Properties

This dialog box is displayed if you right-click a node under a connection in the Versioning navigator and select Properties. It displays properties and property values for the selected object.

Related Topics

- [Using Versioning](#) (page 1-53)

6.253 Subversion: Remove from Subversion

Use this dialog box to begin the process of removing the listed files from the Subversion repository.

After you have clicked **OK**, the listed files will appear on the **Outgoing** tab of the **Pending Changes** window. The files will be removed from the Subversion dialog when you next commit the individual files or the working copy that they are part of.

Related Topics

- [Using Versioning](#) (page 1-53)

6.254 Subversion: Repository Browser

Use this dialog box to select the location of a Subversion repository when using the branching, merging, and switching facilities. Locations in this dialog are shown as directories and objects. The chosen location is ultimately returned from this dialog as a URL.

Repository Connection: If the required location already exists, select it from the browser tree.

To create a new location, navigate to a parent directory, then select the **Create New Remote Directory** icon. This opens a dialog box that will show the location of the parent object (in the form of a URL) and let you name a directory beneath that one that will become the new location.

Related Topics

- [Using Versioning](#) (page 1-53)

6.255 Subversion: Revert Local Changes

Use this dialog box to revert files to their previous state.

If the contents of a file have been changed, the file will be reverted to its base revision. If a file has been added but not yet committed, it will revert to unadded status. If a file is scheduled for removal (in the **Pending Changes** window), it will be added back to the navigator and given its previous status.

Files List: Lists the names and physical locations of the files that will be reverted.

Recursive: Select if you want the revert operation to recurse into child objects of those selected.

Related Topics

- [Using Versioning](#) (page 1-53)

6.256 Subversion: Switch

Use this dialog box to update the current working copy of the specified file from the specified repository and revision.

From URL: Full URL for the repository location associated with the current working copy.

To URL: Full URL for the repository location to use to update the current working copy.

HEAD Revision: Causes the HEAD revision (the latest revision in the repository) to be used for the update operation.

Use Revision: Causes the revision specified in the text box to be used for the update operation. To see a list of revisions from which you can choose, click **List Revisions**.

Related Topics

- [Using Versioning](#) (page 1-53)

6.257 Subversion: Unlock Resources

Use this dialog box to perform a Subversion unlock operation on one or more locked, checked out files (working copies).

Files List: Lists the names and physical locations of the files to be unlocked. You can individually select and deselect files.

Force Unlock: Breaks any existing locks and unlocks the files. Causes the `--force` option to be added to the underlying `svn unlock` command.

Related Topics

- [Using Versioning](#) (page 1-53)

6.258 Subversion: Update Resources

Use this dialog box to incorporate into your local copies changes that other developers have committed to the Subversion repository.

Files List: Lists the names and physical locations of the files that will be updated with content from the Subversion repository.

Use Revision: Updates the files with content from a particular revision within the Subversion repository. Enter the revision number in the adjacent text box. If not selected, the files will be updated from the HEAD revision.

Ignore Externals: Select if you do not want the update operation to apply to external working copies created as the result of externals definition handling. Externals definitions are used to pull data from multiple repositories. See the Subversion documentation for details.

Recursive: Deselect if you do not want the update operation to recurse into child objects of those selected.

Related Topics

- [Using Versioning](#) (page 1-53)

6.259 Subversion: Update Working Copy

Use this dialog box to update individual files with content from the Subversion repository.

Files List: Lists the names and physical locations of the files that will be updated with content from the Subversion repository.

Use Revision: Updates the files with content from a particular revision within the Subversion repository. Enter the revision number in the adjacent text box. If not selected, the files will be updated from the HEAD revision.

Ignore Externals: Select if you do not want the update operation to apply to external working copies created as the result of externals definition handling. Externals definitions are used to pull data from multiple repositories. See the Subversion documentation for details.

Recursive: Deselect if you do not want the update operation to recurse into child objects of those selected.

Related Topics

- [Using Versioning](#) (page 1-53)

6.260 Subversion: Versioning Properties

This dialog box displays general and versioning information about the currently selected file or folder.

Related Topics

- [Using Versioning](#) (page 1-53)

6.261 Third-Party Database Objects

You have requested help about a type of object in the context of a third-party (non-Oracle) database connection. See the documentation for the third-party database for information about database objects as they apply to that product.

6.262 Unable to Connect

This box informs you that SQL Developer is unable to connect to the Internet. The cause might be that the connection information for the specified HTTP proxy server is invalid or the server is not available.

6.263 Unable to Open File

This box informs you that SQL Developer is unable to perform the export operation to the location and file that you specified. The cause might be that you do not have permission to write to that location.

6.264 Uninstall Oracle REST Data Services

This wizard is displayed if you click Tools, then REST Data Services, then Uninstall. Uninstalling Oracle REST Data Services performs the following actions:

- Removes the Oracle REST Data Services schema (ORDS_METADATA)
- Removes the Oracle REST Data Services proxy user (ORDS_PUBLIC_USER)
- Removes Oracle REST Data Services-related database objects (public synonyms and other objects)

If you want to reinstall Oracle REST Data Services, you must first uninstall the existing Oracle REST Data Services; and before you uninstall, you must ensure that Oracle REST Data Services is stopped.

ORDS File Locations: For the Oracle REST Data Services `ords.war` file, specify whether to use the one that is included with SQL Developer or the one that you downloaded from the Oracle Technology Network for uninstalling the Oracle REST Data Services schema.

Database Connection: Specify the database connection information for the database from which to uninstall Oracle REST Data Services: host, port, and the database name (SID) or service name. Additionally, specify the user name and password for a user with SYSDBA privileges to uninstall Oracle REST Data Services (enter SYS credentials).

Database Options: This step of the wizard appears only for a multitenant container database (CDB). Select the CDB and its PDBs from which to uninstall Oracle REST Data Services.

Summary: Click **Finish** to start the process.

6.265 Unit Testing: Action Required

This dialog box is displayed if you are creating a repository but the database user for the repository does not have the necessary privileges.

SQL Developer will issue one or more prompts. For each, click **Yes**, enter the password for the SYS account, and allow the commands to execute.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.266 Unit Testing: Add Category

Use this dialog box to create a lookup category for unit testing.

Related Topics:

[Using Lookups to Simplify Unit Test Creation](#) (page 4-8)

[SQL Developer: Unit Testing](#) (page 4-1)

6.267 Unit Testing: Add Data Type

Use this dialog box to add a data type to a lookup category for unit testing.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.268 Unit Testing: Add Item to Library

Use this dialog box to add a startup, teardown, or validation specification to the unit testing library, depending on which node you have selected in the Unit Test navigator. After you specify a name and click OK, specify a connection, and then complete the definition.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.269 Unit Testing: Add Test Implementation

Use this dialog box to add an implementation for a unit test.

When you create a unit test, a default implementation is created; however, you can specify one or more additional implementations.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.270 Unit Testing: Add Test Suite

Use this dialog box to create a test suite for unit testing.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.271 Unit Testing: Add Tests or Suites to a Suite

Use this dialog box to add one or more unit tests or suites to a test suite. The Tests or Suites tab is displayed, depending on whether you right-clicked the Tests or Suites node in the Unit Test navigator display and selected Add Test or Add Suite.

List of tests or suites: Select one or more tests to be added to the suite.

Run Test Startups: If this option is checked, the startup action defined for each specified test or suite is run when the suite is run.

Run Test Teardowns: If this option is checked, the teardown action defined for each specified test or suite is run when the suite is run.

Any startup and teardown actions for a test are specified when you create the unit test.

Related Topics:

[SQL Developer: Unit Testing](#) (page 4-1)

[Unit Testing: Create Unit Test](#) (page 6-152)

6.272 Unit Testing: Copy or Rename Unit Test

Use this dialog box to copy or rename a unit test.

If you specified Copy, a copy of the selected unit test is created and is given the name you specify.

If you specified Rename, the name of the selected unit test is changed to the name you specify.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.273 Unit Testing: Create Unit Test

Use this wizard to create a unit test. To go from one step to the next, click Next; to go back to the previous step, click Back.

Related Topics:

[SQL Developer: Unit Testing](#) (page 4-1)

6.273.1 Select Operation

Select a database connection on the right, then use the hierarchy tree to select an object to be tested. For example, to test a procedure, expand the Procedures node and select the desired procedure.

6.273.2 Specify Test Name

Test Name: Name for the unit test. Can be the same as the name of the database object (for example, procedure or function) to be tested.

Create with single dummy representation: Creates a single (sometimes called "one-off") test case for which you must specify the input parameter values when you run the test.

Select/create implementation using lookup values: Generates multiple test cases using sets of input parameter values that you will specify.

6.273.3 Specify Startup

Specify the action to perform at the start of the test, before any of the actual test operations are performed. Reasons for using a startup action might include the following:

- You want to create copies of the entire table or specific rows that will be modified by the test, so that you can restore the original values later during teardown.
- You want to perform some special operations before the test is run.

None: Perform no startup action.

Row Copy: Copy rows from a specified table. You will be asked to specify the table and q query for selecting the rows.

Table Copy: Copy a specified table. You will be asked to specify the table. For example, you might want to copy the EMPLOYEES table to a table named EMPLOYEES_ORIGINAL.

User PL/SQL Code: Run a script. You will be asked to specify the script file.

6.273.4 Specify Parameters

If you are creating a single test case, specify the parameter values in the Input column for each input parameter. For example, for the example `award_bonus` procedure, you might test the case where employee ID 1001 sold 5000 dollars work of goods or services (EMP_ID Input: 1001, SALES_AMT Input: 5000).

If you are creating multiple test cases, specify the information for each parameter for each test case. For example, to create a test with 5 test cases for the example `award_bonus` procedure, you need to enter 10 rows on this page (2 input parameters times 5 test cases).

6.273.5 Specify Validations

Specify one or more validation actions to perform after the test case is run, to check whether the desired result was returned. For example, if a test case for a procedure

was supposed to increase the salary for employee number 1001 to \$2000, specify a Query returning row(s) validation of `SELECT salary FROM employees WHERE employee_id = 1001 AND salary = 2000`.

For all validation action options except None, you will be prompted to specify the required information.

To add a validation action for a test case, click the Add (+) icon; to delete a validation action for a test case, select it and click the Remove (X) icon.

None: Perform no validation action.

Boolean function: Validation succeeds if Boolean TRUE is returned by the specified function.

Query returning no row(s): Validation succeeds if the specified query returns no rows.

Query returning row(s): Validation succeeds if the specified query returns one or more rows that satisfy the specified query.

Result Set Compare: Validation succeeds if the results returned by the specified query match the values that you specify.

Table Compare: Validation succeeds if the target table (the table modified by the test case) is identical to the specified table.

User PL/SQL Code: Validation succeeds if the specified script runs successfully.

6.273.6 Specify Teardown

Specify the action to perform after the validation action (or the test case in no validation action was specified) is finished. Reasons for using a teardown action might include the following:

- You want to restore a table to its original values, effectively "rolling back" any changes that were made by the test case.
- You want to perform some special operations after the test is run.

None: Perform no teardown action.

Table Drop: Drop (delete) the specified table.

Table Restore: Replace all rows in the specified target table (the table modified by the test case) with the rows in the specified source table. For example, you might want to replace the contents of EMPLOYEES (target) table with a source table named EMPLOYEES_ORIGINAL.

User PL/SQL Code: Run a script. You will be asked to specify the script file.

6.273.7 Summary

Displays the selected options. To make any changes, click Back. To perform the operation, click Finish.'

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.274 Unit Testing: Manage Users

Use this dialog box to add or remove database users as unit test users (Users tab) or administrators (Administrators tab).

On each tab, use the icons to move selected database users (> or <) or all listed database users (>> or <<) into or out of the column on the right, which identifies the unit testing users or administrators.

With Admin: If this option is checked, the users in the column on the right can grant the specified unit test role (User or Administrator) to other database users.

Related Topics:

[SQL Developer: Unit Testing](#) (page 4-1)

[Managing Repository Users and Administrators](#) (page 4-5)

6.275 Unit Testing: Rename Test Implementation

Use this dialog box to rename an implementation for a unit test.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.276 Unit Testing: Result of Operation

This informational box indicates either that the operation was successful or that an error occurred. If an error occurred, a brief explanation is included.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.277 Unit Testing: Synchronize Unit Test

Use this wizard to synchronize a unit test, that is, to update the selected unit test to reflect changes in the operation to be tested. For example, you might want the unit test to use a new procedure (perhaps owned by a different database user) instead of the procedure for which the unit test was originally defined. (The SQL Developer unit testing feature is explained in [SQL Developer: Unit Testing](#) (page 4-1).) To go from one step to the next, click Next; to go back to the previous step, click Back.

Select Operation

Select the database connection (associated with the owner of the object) on the right, then use the hierarchy tree to select an object to be used as the new operation. For example, to specify a different procedure, expand the Procedures node and select the desired procedure.

Current Target Operation: The operation currently associated with this unit test.

Select New Target Operation: Use the hierarchy tree to select the new operation to be associated with this unit test.

Reset: Discards any unsaved specifications and restores the current definition.

Match Parameters by Name/Position: Specify whether to match the new operation's parameters with the current operation's parameters by using the names or positions of the parameters in the format (signature) of the object.

Connection selector: Select the database connection to be used for displaying the hierarchy of objects from which to select the new operation. For example, if the new procedure is in a different database schema, select a connection associated with the database user for that schema.

New Target Operation: Displays the result of your selection.

Summary

Displays information about the current operation, the new operation, and the changes that will be made to the unit test. To make any changes, click Back. To perform the operation, click Finish.

Related Topics

- [SQL Developer: Unit Testing](#) (page 4-1)

6.278 Unplug Pluggable Database

Use this dialog box to unplug a PDB from a CDB. When you unplug a PDB, Oracle stores metadata for the PDB in an XML file. You can use this XML file to subsequently plug the PDB into a CDB.

Database Name: Name of the PDB.

XML File Name: The full path of the XML file in which to store the metadata for the unplugged PDB.

Related Topics

- [Multitenant Container Database \(CDB\)](#) (page 1-19)

6.279 Unsupported Database Version

This box is displayed if you try to create a connection to a database release that is not supported by SQL Developer, such as Oracle Database release 8.1. For information about database releases supported by SQL Developer, see *Oracle SQL Developer Installation Guide*.

6.280 Validate Geometry

This dialog box is displayed if you select a table and specify Validate Geometry Using Tolerance or Validate Geometry Using Dimension Information. Internally, SQL Developer calls the SDO_GEOM.VALIDATE_LAYER_WITH_CONTEXT procedure, which examines the specified geometry column to determine if the stored geometries follow the defined rules for geometry objects, and returns context information about any invalid geometries. For information about the specific consistency checks performed, see the documentation for the SDO_GEOM.VALIDATE_LAYER_WITH_CONTEXT procedure in *Oracle Spatial and Graph Developer's Guide*.

Table Name: Name of the feature table that has a column of type SDO_GEOMETRY.

Tolerance (for Validate Geometry Using Tolerance): Tolerance value to be used for validating the geometries. For example, you might try larger values to see if more

geometries would be invalid, or smaller values to see if fewer geometries would be invalid.

Spatial Column: Name of the column of type SDO_GEOMETRY in the specified feature table.

Identification Columns: Other (non-spatial) columns in the table.

Validation Results: For any invalid geometries, identifies the ROWID value of the row containing the invalid geometry, and an Oracle error message number and the context of the error (the coordinate, edge, or ring that causes the geometry to be invalid). You can then look up the error message for more information about the cause of the failure.

Related Topics

- [Spatial Support in SQL Developer](#) (page 1-71)

6.281 View Whole Value

This dialog box is displayed if, during debugging, you right-click an item (node) of type string, byte array, char array, short array, int array, or long array in the Data tab display and select View Whole Value.

View as String: Displays the item's value as a text string, as opposed to hexadecimal values.

Wrap Text: Enables the wrapping of long lines in the box in which the value is displayed.

6.282 Windows

This dialog box is displayed if you right-click the tab for a window in the display area of the SQL Developer main window.

Windows: A list of the windows in the display area.

Activate: Makes active (switches focus to) the selected window.

Close: Closes the selected window.

6.283 XMLType Column Properties

This dialog box is displayed if you define a table column as type SYS.XMLTYPE and click the XMLType Properties button.

Store As: **Binary XML** (data stored in compact, post-parse, XML schema-aware format; can be used with XML data that is not based on an XML schema), **CLOB** (data stored in Character Large Object instances), or **Object Relational** (data stored as a set of objects). Note that effective with Oracle Database Release 12.1, Store As CLOB for XMLType is deprecated.

XML Schema Owner: Owner of the XML schema.

XML Schema: Name of the XML Schema. For example: `http://xmlns.oracle.com/xdm/XDBResource.xsd`

Element: Select the desired element from the list of elements defined in the XML schema.

Allow Any Schema (only if Store As is Binary XML): Encodes all XML schema-based documents using an encoding based on the XML schema referenced by the document.

Allow Non-Schema (only if Store As is Binary XML): Encodes all XML schema-based documents using an encoding based on the referenced XML schema; encodes all non-schema-based documents using the non-schema-based encoding.

SQL Developer: TimesTen Support

If you connect to an Oracle TimesTen In-Memory Database, you can perform many TimesTen-specific operations within SQL Developer.

This topic adapts selected material from *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*, which is in the SQL Developer documentation library.

- [Create a New TimesTen Connection](#) (page 7-2)
- [Set the Cache Administrator User Name and Password](#) (page 7-5)
- [Start and Stop the Cache Agent](#) (page 7-6)
- [Create a Cache Group](#) (page 7-6)
- [Drop a Cache Group](#) (page 7-13)
- [Load or Refresh a Cache Group](#) (page 7-13)
- [Flush a Cache Group](#) (page 7-14)
- [Unload a Cache Group](#) (page 7-14)
- [Alter the Automatic Refresh Attributes of a Cache Group](#) (page 7-15)
- [Specify and Display LOB Data Types](#) (page 7-15)
- [Specify the INLINE Attribute for Columns](#) (page 7-15)
- [Compute Table Size](#) (page 7-16)
- [Use a REF CURSOR as an OUTPUT Parameter](#) (page 7-17)
- [Create an Index on a TimesTen Table](#) (page 7-17)
- [Create a Primary Key Constraint](#) (page 7-18)
- [Change the Attributes of Primary Key](#) (page 7-18)
- [Create a Foreign Key Constraint](#) (page 7-19)
- [Drop a Constraint](#) (page 7-20)
- [Add an Aging Policy to a Table](#) (page 7-20)
- [Drop an Aging Policy from a Table](#) (page 7-21)
- [Change the Aging State for the Aging Policy of a Table](#) (page 7-21)
- [Change the Memory Usage Thresholds and LRU Aging Cycle](#) (page 7-21)

- [Change the Lifetime and Cycle for the Time-Based Aging Policy of a Table](#) (page 7-21)
- [Grant Object Privileges to a TimesTen Database Element](#) (page 7-22)
- [Revoke Object Privileges from a TimesTen Database Element](#) (page 7-22)
- [Load Data into a Table](#) (page 7-23)
- [Start and Stop the TimesTen Replication Agent](#) (page 7-23)
- [Update the Table and Column Statistics Used By the TimesTen Query Optimizer](#) (page 7-23)
- [Export the Table Statistics Used By the TimesTen Query Optimizer](#) (page 7-25)
- [Create a TimesTen Performance Metrics Snapshot](#) (page 7-25)
- [Drop a TimesTen Performance Metrics Snapshot](#) (page 7-26)
- [Edit the TimesTen Performance Metrics Configuration](#) (page 7-27)
- [Create a TimesTen Performance Metrics Report](#) (page 7-27)
- [Capture Data for the TimesTen Index Advisor at the Connection Level](#) (page 7-29)
- [Capture Data for the TimesTen Index Advisor at the Database Level](#) (page 7-33)
- [View the TimesTen Index Advisor Capture State](#) (page 7-35)
- [Load Data from an Oracle Database into a TimesTen Database](#) (page 7-36)
- [Load Data from an Oracle Database Table into an Existing TimesTen Table](#) (page 7-39)

7.1 Create a New TimesTen Connection

To create a named connection to a TimesTen database, make sure you are on the main SQL Developer page.

1. Right-click the **Connections** node and select **New Connection**.
The **New/Select Database Connection** dialog appears.
2. Select the **TimesTen** tab.
You are ready to enter information for the connection.
3. In the **Connection Name** field, enter the connection name.
4. In the **Username** field, enter the TimesTen user. If you use the TimesTen database to cache data from an Oracle database, enter the TimesTen cache manager user name. If you use the TimesTen database to load data from an Oracle database, enter the TimesTen user name that has **SELECT** privileges on the Oracle database tables that you want to load.
5. In the **Password** field, enter the password for the TimesTen user.

To save the password, choose **Save Password**. If you save the password, you will not be prompted for the password on subsequent connections to the TimesTen database.

6. To allow TimesTen to issue an implicit commit after DML statements, select the **Autocommit** check box. By default, the **Autocommit** check box is selected.
7. Select a connection type from the **Connection Type** drop-down list. The following types of connections are available:
 - DSN
 - Easy Client/Server
 - Advanced

You are now ready to define your new connection.

Related Topics:

[DSN](#) (page 7-3)

[Easy Client/Server](#) (page 7-4)

[Advanced](#) (page 7-5)

7.1.1 DSN

The DSN connection type allows you to connect to databases by specifying a DSN. The DSN must exist in your `.ODBC.INI` on UNIX and the ODBC Data Source Administrator on Windows.

1. Select a predefined DSN that references the TimesTen database from the **DSN** drop-down list. You can select either a Data Manager DSN to establish a direct connection or a client DSN to establish a client/server connection.
2. In the **Connection Attributes** field, enter the attributes for the DSN. The **Connection Attributes** field is not required. Each attribute setting must be separated by a semicolon. If you do not specify any connection attributes, TimesTen uses the connection attributes specified in the DSN.
3. To cache data from an Oracle database or load data from an Oracle database into a TimesTen database, select the **Enable Cache/Load tables from Oracle database** check box.
4. In the **Oracle Database Password** field, enter the Oracle database password for the Oracle database user. The Oracle database user must be the same user as the TimesTen user you specified in the **Username** field. This text field is only available if you selected the **Enable Cache/Load tables from Oracle database** check box.
5. Click **Save** to save the connection name under the **Connections** node.
Your connection is saved. You can establish connections to the database using the settings in the named connection without having to specify the values in the fields of the **New/Select Database Connection** dialog box again.

6. Click **Connect**.

SQL Developer is connected to your TimesTen database.

See Also:

["Specifying Data Source Names to identify TimesTen databases"](#) in the *Oracle TimesTen In-Memory Database Operations Guide*

["Connection Attributes"](#) in the *Oracle TimesTen In-Memory Database Reference Database Connections* (page 1-28)

7.1.2 Easy Client/Server

The easy client/server connection type allows you to connect to local and remote databases without having to define a DSN.

1. In the **Server Host Name** field, enter the name or the IP address of the TimesTen server host.
2. In the **Server DSN** field, enter the DSN of the TimesTen server.
3. In the **Server Port** field, enter the TCP port of the TimesTen server. You can use the `ttStatus` utility to find the TCP port of your TimesTen server.

In the following example, the server port number is 53396. This example is part of a sample output from the `ttStatus` utility.

```
TimesTen server pid 24379 started on port 53396
```

For more information about the `ttStatus` utility, see "ttStatus" in the *Oracle TimesTen In-Memory Database Reference*.

4. To cache data from an Oracle database or load data from an Oracle database into a TimesTen database, select the **Enable Cache/Load tables from Oracle database** check box.
5. In the **Oracle Database Password** field, enter the Oracle database password for the Oracle database user. The Oracle database user must be the same user as the TimesTen user you specified in the **Username** field. This text field is only available if you selected the **Enable Cache/Load tables from Oracle database** check box.
6. In the **Oracle Net Service Name** field, enter the Oracle database net service name for the Oracle database. This text field is only available if you selected the **Enable Cache/Load tables from Oracle database** check box.
7. Click **Save** to save the connection name under the **Connections** node.

Your connection is saved. You can establish connections to the database using the settings in the named connection without having to specify the values in the fields of the **New/Select Database Connection** dialog box again.

8. Click **Connect**.

SQL Developer is connected to your TimesTen database.

Related Topics:

[Database Connections](#) (page 1-28)

7.1.3 Advanced

The advanced connection type allows you to define a connection string to connect to your TimesTen database.

1. Select a driver type for the TimesTen database from the **Driver Type** option. Select either **Client/server** to establish a client/server connection or **Direct** to establish a direct connection.
2. In the **Connection string** field, enter a connection string to define the DSN and set attributes for the DSN. If you choose Client/Server as the Driver Type specify the server host, server port, and DSN. Each attribute setting must be separated by a semicolon. Do not delimit the connection string with a pair of double quotation marks.
3. To cache data from an Oracle database or load data from an Oracle database into a TimesTen database, select the **Enable Cache/Load tables from Oracle database** check box.
4. In the **Oracle Database Password** field, enter the Oracle database password for the Oracle database user. The Oracle database user must be the same user as the TimesTen user you specified in the **Username** field. This text field is only available if you selected the **Enable Cache/Load tables from Oracle database** check box.
5. Click **Save** to save the connection name under the **Connections** node.
Your connection is saved. You can establish connections to the database using the settings in the named connection without having to specify the values in the fields of the **New/Select Database Connection** dialog box again.
6. Click **Connect**.
SQL Developer is connected to your TimesTen database.

See Also:

["Connecting to a database using a connection string"](#) in the *Oracle TimesTen In-Memory Database Operations Guide*

[Database Connections](#) (page 1-28)

7.2 Set the Cache Administrator User Name and Password

You must create a cache administration user in the Oracle database. This user creates, owns and maintains Oracle objects that store information used to manage a specific cache grid. This user also enforces predefined behaviors of cache groups.

After you create the cache administration user in the Oracle database, you must set the name and password for the user in the TimesTen database.

To set the cache administration user name and password, right-click the node of the connection name for the TimesTen database and select **Set Cache Administration User**.

The **Setting the cache administration user** dialog appears. For the **Prompts** tab, in the **Cache administration user ID** field, type the cache administration user name. In the

Cache administration user password field, type the Oracle password for the cache administration user.

Click **Apply**.

You must set the cache administration user name and password for each TimesTen database that caches Oracle. If you drop and re-create the TimesTen database or drop and re-create the cache administration user in the Oracle database, then the cache administration user name and password must be reset in the TimesTen database.

See Also:

["Set the cache administration user name and password"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.3 Start and Stop the Cache Agent

After you set the cache administration user name and password, create a cache grid and associate the TimesTen database with the cache grid, you must start the TimesTen cache agent process. The cache agent processes cache operations such as loading a cache group and automatic refresh.

To start the cache agent, right-click the node of the connection name for the TimesTen database and select **Start/Stop Cache Agent**.

In the **Starting/stopping the cache agent** dialog box, select **Start agent** within the **Prompts** tab. Then click **Apply** to start the cache agent.

To stop a running cache agent process, right-click the node of the connection name for the TimesTen database and select **Start/Stop Cache Agent** in an identical fashion as was done when starting the cache agent.

In the **Starting/stopping the cache agent** dialog box, select **Stop agent** within the **Prompts** tab. Then click **Apply** to stop the cache agent.

See Also:

["Managing the cache agent"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.4 Create a Cache Group

A cache group defines the set of Oracle data to cache in the TimesTen database. When you create a cache group, cache tables are created in the TimesTen database that correspond to the Oracle tables being cached.

To create a cache group, right-click the **Cache Groups** node under the node of the connection name for the TimesTen database and select **New Cache Group**.

In the **Create cache group** dialog, specify a name for the cache group in the **Name** field within the **Cache Group Attributes** tab. For **Type**, choose the desired cache group type (read-only, asynchronous writethrough, synchronous writethrough or user managed).

Cache groups are categorized as either explicitly loaded or dynamic. If you are creating a dynamic cache group, then choose **Dynamic** in the **Type** section.

In an explicitly loaded cache group, data is loaded manually into its cache tables from the corresponding Oracle tables using a load or refresh operation, or automatically using an automatic refresh operation. In a dynamic cache group, data is loaded into its cache tables on demand from the corresponding Oracle tables using a dynamic load operation, or manually using a load operation.

A read-only, asynchronous writethrough (AWT) or synchronous writethrough (SWT) cache group can be defined as a dynamic cache group. A user managed cache group can be defined as a dynamic cache group if it does not contain all of the following characteristics:

- Automatic refresh is defined
- At least one of its cache tables are updatable
- Committed updates on one or more of its updatable cache tables are propagated to the corresponding Oracle tables

Cache groups are classified as either local or global. If you are creating a global cache group, then select the **Global** in the **Type** section.

In a local cache group, data in its cache tables are not shared across TimesTen databases even if the databases are members of the same cache grid. In a global cache group, data in its cache tables are shared across TimesTen databases that are members of the same cache grid.

Only an AWT cache group can be defined as global cache group.

See also the following related topics:

- [Defining a cache group and cache table attributes](#) (page 7-7)
- [Defining the root table](#) (page 7-8)
- [Defining the child tables](#) (page 7-11)

See Also:

["Loading data into a cache group: Explicitly loaded and dynamic cache groups"](#) and ["Dynamic cache groups"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

["Sharing data across a cache grid: Local and global cache groups"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.4.1 Defining a cache group and cache table attributes

Within the **Cache Group Attributes** tab of the **Create cache group** dialog, if you are creating a user managed cache group and you select **Autorefresh** within the **Table Attributes** section to enable automatic refresh when either **All tables readonly** or **Propagate all tables** is selected, or you are creating a read-only cache group, then for the following automatic refresh characteristics:

- Set the automatic refresh mode in the **Mode** field by choosing either **Full** or **Incremental**. The default automatic refresh mode is Incremental.
- Set the automatic refresh interval in the **Refresh Interval** field by specifying a numeric value followed by a unit of time in minutes, seconds or milliseconds. The default automatic refresh interval is 5 minutes.

- Set the automatic refresh state in the **Refresh State** field by choosing either **Paused**, **On** or **Off**. The default automatic refresh state is **Paused**.

If you are creating a user managed cache group, then from the **Table Attributes** section in the **Cache Group Attributes** tab within the **Create cache group** dialog, choose one:

- **All tables readonly**: All cache tables are read-only
- **Propagate all tables**: All cache tables are updatable and committed updates on the cache tables are propagated to the corresponding Oracle tables
- **Varies from table to table**: Some cache tables are read-only and other cache tables are updatable where committed updates on the cache tables may or may not be propagated to the corresponding Oracle tables

See Also:

["AUTOREFRESH cache group attribute"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

["Read-only cache group"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

["User managed cache group"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.4.2 Defining the root table

Click the **Tables** tab in the **Create cache group** dialog. Right-click in the **Tables** panel and select **Add root table**.

The **Select the root table** dialog appears. The list of Oracle tables and Oracle synonyms that you can choose to designate as the root table for your cache group is displayed. Oracle tables that have primary keys and Oracle tables that have unique non-null indexes are available for selection.

Select the table and click **OK**.

To show only Oracle tables owned by the current connection user, choose **Show only my tables**.

See Also:

["Caching Oracle Synonyms"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.4.2.1 Selecting the columns to cache

In the **Columns** tab within the **Tables** tab of the **Create cache group** dialog, select the check boxes under the **Inline** field to specify the **INLINE** attribute for the column. You can specify the **INLINE** attribute for columns of type **VARCHAR2**, **NVARCHAR2**, and **VARBINARY**.

Select the check boxes under the **Allow null values** field for each cached column to accept null values. By default, the nullability of a cache table's column is identical to the nullability of the corresponding Oracle table's column.

Select the check boxes under the **Cached** field for each column to cache. By default, all columns of the table are selected and are cached. The Oracle table's unique identifying columns (primary key, unique index or unique constraint) must be included in the TimesTen cache table and constitutes the cache table's primary key.

To select all columns or deselect all columns, except for the Oracle table's unique identifying columns, to accept null values or to cache, select or deselect the check boxes that appear to the left of the **Allow null values** and **Cached** fields, respectively.

Some Oracle data types have multiple compatible TimesTen data types. For columns in the Oracle table whose data type can map multiple TimesTen data types, a menu appears under the **Data type in TimesTen** field for those columns. Select the desired data type for the column in the cache table. The default compatible TimesTen data type for the Oracle data type appears in italics.

For updatable cache tables, you can specify a default column value under the **Default Value** field for each cached column to define a default value. By default, the default value of a cache table's column is identical to the default value of the corresponding Oracle table's column.

See Also:

["Cache groups and cache tables"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

["Mappings between Oracle and TimesTen data types"](#) in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.4.2.2 Specifying the table attributes

By default, TimesTen creates a range index on the cache table's primary key columns. To create a hash index instead of a range index on the primary key columns, in the **Primary Key** section of the **Table Attributes** tab within the **Tables** tab of **Create cache group**, choose **Use hash**. You must size the hash index by specifying either the maximum number of rows the underlying cache table is anticipated to have or the number of pages to allocate to the index in the **Hash sizing policy** menu and the adjacent field after the colon (:) which requires a numeric value. The default selection is **Current table row count**.

For more information about the index that is automatically created on a cache table's primary key columns, see "CREATE CACHE GROUP" in the *Oracle TimesTen In-Memory Database SQL Reference*.

For more information about properly sizing a hash index, see "CREATE TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference*.

By default, the primary key columns of a cache table are identical to the primary key columns of the corresponding Oracle database table. You can define a different column or set of columns as the primary key for the cache table.

In the **Manual table management** section of the **Table Attributes** tab within the **Tables** tab of the **Create cache group** dialog, click **Redefine primary key**.

In the **Define Primary Key(s)** dialog, the non-nullable unique columns of the corresponding Oracle table is shown in the **Non-null unique columns** list. Select the columns for the primary key of the cache table and then choose >> to move those columns into the **Primary key(s)** list. The order of the columns in a composite primary key is significant if a range index is to be created on the columns. Click **OK** to define the alternate primary key for the cache table.

If the primary key of the cache table consists of a different set of columns than the primary key of the corresponding Oracle table, then the tables may become unsynchronized when committed updates are made on the cache table or the corresponding Oracle table. Unique constraint violations can occur when:

- Committed updates on the cache table are propagated to the corresponding Oracle table
- Committed updates on the Oracle table are refreshed to the corresponding cache table

You can specify an optional WHERE clause for the cache table in the **WHERE clause** field of the **Table Attributes** tab within the **Tables** tab of **Create cache group**. A WHERE clause is only allowed for table definitions in a read-only cache group or a user managed cache group that has automatic refresh defined.

For more information about specifying a WHERE clause in the table definition of a cache group, see "Using a WHERE clause" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

If the cache table is in a user managed cache group that has selected the **Varies from table to table** in the **Table attributes** section within the **Cache Group Attributes** tab of the **Create cache group** dialog, then you must specify the table attribute for the cache table.

In the **Propagation Type** section of the **Table Attributes** tab within the **Tables** tab of **Create cache group**, choose:

- **Readonly** to define the cache table as read-only
- **Propagate** to define the cache table as updatable with updates on the table to be propagated to the corresponding Oracle table
- **Neither** to define the cache table as updatable with updates on the table not propagated to the corresponding Oracle table

By default, cache tables in a user managed cache group are updatable with updates on the table not propagated to the corresponding Oracle table.

7.4.2.3 Specifying an aging policy

You can specify an aging policy for the cache table in the **Aging** tab of the **Create cache group** dialog. An aging policy is explicitly defined on the root table of a cache group and it implicitly applies to all the child tables.

Choose:

- **Usage-based** to define a least recently used (LRU) aging policy
- **Time-based** to define a time-based aging policy
- **No aging** to not define an aging policy

An LRU aging policy can be defined for tables in an explicitly loaded AWT, SWT or user managed cache group that does not have automatic refresh defined. The default aging state is set to off for global explicitly loaded cache groups. For dynamic cache groups that do not have automatic refresh defined, the default aging state is set to on and the type of aging is set to usage-based.

The TimesTen database memory usage thresholds that determine when data starts and stops being deleted from tables that are subject to LRU aging are shown when the

Usage-based is selected. The default usage threshold that determines when data starts being deleted from the tables is 90%. The default usage threshold that determines when data stops being deleted from the tables is 80%. The default LRU aging cycle is 1 minute.

A time-based aging policy can be defined on a cache table that contains a non-nullable DATE or TIMESTAMP column. Columns in the table of these data types appear in the **Column** within the **Aging** tab of the **Create cache group** dialog when **Time-based** is selected. Select the column to store the timestamp value indicating when each row was inserted or most recently updated.

In the **Life-time** field, indicate the length of time in which data that has not been updated is to be kept in the table by specifying a numeric value followed by a unit of time in minutes, hours or days from the drop-down menu list.

In the **Cycle** field, indicate the frequency in which data is to be aged out of the table by specifying a numeric value followed by a unit of time in minutes, hours or days in the drop-down menu list. The default time-based aging cycle is 5 minutes.

Specify an aging state by choosing **on** or **off**.

See Also:

"LRU aging" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*
[Change the Memory Usage Thresholds and LRU Aging Cycle](#) (page 7-21)

"Time-based aging" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.4.2.4 Verifying the root table definition

After you have defined the tables and columns to cache, the type of index to create on the primary key columns, the table attributes including an optional WHERE clause, and an optional aging policy for the root table, you can view the CREATE CACHE GROUP statement used to create the cache group and its cache table in the **DDL** tab of the **Create cache group** dialog.

Click **Apply** to create the cache group and its cache table if you are not going to add any child tables to the cache group.

For information about the syntax and semantics of the CREATE CACHE GROUP statement, see "CREATE CACHE GROUP" in the *Oracle TimesTen In-Memory Database SQL Reference*.

7.4.3 Defining the child tables

Within the **Tables** tab of the **Create cache group** dialog, right-click in the **Tables** panel and select **Add child tables**.

The **Select one or more child tables** dialog appears. This dialog displays the list of Oracle tables and Oracle synonyms that you can choose to designate as the child table(s) for your cache group. Oracle tables that have primary keys and Oracle tables that have unique non-null indexes are available for selection.

Select the tables to designate as the child tables for your cache group and click **OK**.

For more information on caching Oracle synonyms, see "Caching Oracle Synonyms" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

To select multiple child tables, press and hold the CTRL key, and click the desired tables. The child tables appear under their appropriate parent table. You can also drag and drop a particular child table so that it appears under the desired parent table.

To show only Oracle tables owned by the current connection user, choose **Show only my tables**.

7.4.3.1 Selecting the columns to cache

Click a child table and in the **Columns** tab within the **Tables** tab of the **Create cache group** dialog, select the check boxes under the **Allow null values** field for each cached column to accept null values. By default, the nullability of a cache table's column is identical to the nullability of the corresponding Oracle table's column.

Select the check boxes under the **Cached** field for each column to cache. By default, all columns of the table are selected and are cached. The Oracle table's unique identifying columns must be included in the TimesTen cache table and constitute the cache table's primary key. The Oracle table's foreign key columns that reference its parent table must also be included in the cache table.

To select all columns or deselect all columns, except for the Oracle table's unique identifying columns and the foreign key columns that reference its parent table, to accept null values or to cache, select or deselect the check boxes that appear to the left of the **Allow null values** and **Cache** fields.

For columns in the Oracle table whose data type can map to multiple TimesTen data types, a drop-down select list appears under the **Data type in TimesTen** field for those columns. Select the desired data type for the column in the child cache table.

For updatable cache tables, you can specify a default column value under the **Default Value** field for each cached column in the child cache table that defines a default value.

7.4.3.2 Specifying the table attributes

In the **Table Attributes** tab within the **Tables** tab of the **Create cache group** dialog, select whether to create a range index or a hash index on the child cache table's primary key columns in a similar fashion as was specified for the root table. You can define a different column or set of columns as the primary key for the cache table than the columns of the corresponding Oracle table in a similar fashion as was specified for the root table. You can also specify an optional WHERE clause for the cache table in a similar fashion as was specified for the root table. A WHERE clause is only allowed for table definitions in a read-only cache group or a user managed cache group that has automatic refresh defined.

If you would like the foreign key columns of a child table to have cascade delete enabled such that when rows containing referenced key values are deleted from its parent table, rows in the child table with dependent foreign key values are automatically deleted, then select **Cascade delete foreign key**.

By default, the foreign key columns of a child cache table are identical to the foreign key columns of the corresponding Oracle table. You can define a different column or set of columns as the foreign key for the cache table.

In the **Manual table management** section of the **Table Attributes** tab within the **Tables** tab of the **Create cache group** dialog, select **Redefine foreign key**.

In the **Define Foreign Key(s)** dialog, the columns in the child table whose data types are compatible with the primary key columns of the parent table are shown in the

Foreign Key(s) drop-down select list. Select the columns for the foreign key of the cache table. Click **OK** to define the alternate foreign key for the cache table.

For more information about the hierarchy and relationship of cache tables in a multiple-table cache group, see "Multiple-table cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

If the child cache table is in a user managed cache group that has selected **Varies from table to table** in the **Table attributes** section within the **Cache Group Attributes** tab of the **Create cache group** dialog, then you must specify the table attribute for the cache table in a similar fashion as was specified for the root table. If the cache table is updatable, then specify whether updates on the table are to be propagated or not propagated to the corresponding Oracle table.

7.4.3.3 Completing the definitions of the child tables

Define the columns to cache, the type of index to create on the primary key columns, the cascade delete setting for the foreign key columns, and the table attributes including an optional WHERE clause for any remaining child tables.

7.4.3.4 Verifying the cache group definition

In the **DDL** tab of the **Create cache group** dialog, you can view the CREATE CACHE GROUP statement used to create the cache group and its cache tables.

Click **Apply** to create the cache group and its cache tables.

If you created an AWT cache group, then you must start the replication agent for committed updates on its cache tables to be asynchronously propagated to the corresponding Oracle tables. For information about how to start the replication agent, see [Start and Stop the TimesTen Replication Agent](#) (page 7-23).

If you created a global cache group, then you must attach the TimesTen database to the cache grid that it is associated with before you can perform operations on the cache group or on its cache tables. This task must be performed either outside of SQL Developer using a command-line interface such as the *ttIsq1* utility, or within SQL Developer using SQL Worksheet. See "Attaching a TimesTen database to a cache grid" in the *Oracle TimesTen Application-Tier Database Cache User's Guide* for more information.

7.5 Drop a Cache Group

When you drop a cache group, its cache tables are also dropped.

Under the **Cache Groups** node within the node of the connection name for the TimesTen database, right-click the name of the cache group to drop and select **Drop**.

Click **Apply** to drop the cache group.

For more information about dropping a cache group, see "Dropping a cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.6 Load or Refresh a Cache Group

Click the + to the left of the **Cache Groups** node to view the list of cache groups. Next right-click the name of the cache group to load or refresh. Select **Load** to load the cache group or **Refresh** to refresh the cache group.

The cache agent must be running to load or refresh a cache group. For information about how to start the cache agent, see [Start and Stop the Cache Agent](#) (page 7-6).

For a read-only cache group or a user managed cache group that has automatic refresh defined, you must set the automatic refresh state to Paused before you can load or refresh the cache group. For information about changing the automatic refresh state of a cache group, see [Alter the Automatic Refresh Attributes of a Cache Group](#) (page 7-15).

In the **Prompts** tab of the **Load** or **Refresh** dialog, specify a numeric value in the **Commit every n rows** field to indicate the frequency, based on the number of rows that are loaded or refreshed into the cache group, at which a commit is issued during the load or refresh operation. By default, a transaction is committed after every 256 rows are loaded or refreshed into the cache group.

In the **Number of threads to run in parallel** field within the **Prompts** tab of the **Load** or **Refresh** dialog box, specify the number of cache agent process threads to spawn and use for the load or refresh operation. Parallel processing of the load or refresh operation can potentially improve the operation's throughput. The numeric value specified in this field cannot exceed 10 and should not exceed the number of CPUs on the system that can be dedicated to the parallel load or refresh operation. The default value of 1 indicates that the load or refresh operation is processed serially. For more information about parallel loading or refreshing of cache groups, see "Improving the performance of loading or refreshing a large number of cache instances" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

You can specify an optional WHERE clause for the load or refresh operation in the **Where clause** field within the **Prompts** tab of the **Load** or **Refresh** dialog. You can use a WHERE clause for a load or refresh operation on an AWT, SWT or user managed cache group that does not have automatic refresh defined. For more information about specifying a WHERE clause in a load or refresh operation, see "Loading and refreshing a cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

Click **Apply** to load or refresh the cache group.

For more information about loading or refreshing a cache group, see "Loading and refreshing a cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.7 Flush a Cache Group

Only user managed cache groups containing cache tables that are not all read-only, or not all updatable where updates on the tables are propagated to the corresponding Oracle tables can be flushed.

Click the + to the left of the **Cache Groups** node to view the list of cache groups. Next right-click the name of the cache group that you to flush. Select **Flush** to flush the cache group.

In the **Where clause** field within the **Prompts** tab of the **Flush** dialog, specify an optional WHERE clause for the flush operation.

Click **Apply** to flush the cache group.

For more information about flushing a cache group, see "Flushing a user managed cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.8 Unload a Cache Group

Click the + to the left of the **Cache Groups** node to view the list of cache groups. Next right-click the name of the cache group to unload. Select **Unload** to unload the cache group.

In the **Where clause** field within the **Prompts** tab of the **Unload** dialog, specify an optional WHERE clause for the unload operation. You can use a WHERE clause for an unload operation on an AWT, SWT or user managed cache group that does not have automatic refresh defined.

Click **Apply** to unload the cache group.

For more information about unloading a cache group, see "Unloading a cache group" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.9 Alter the Automatic Refresh Attributes of a Cache Group

You can alter the automatic refresh attributes of a read-only cache group or a user managed cache group that has automatic refresh defined.

Under the **Cache Groups** node, right-click the name of the cache group and select **Autorefresh**.

In the **Prompts** tab of the **Autorefresh** dialog, you can change any of the following automatic refresh attributes:

- To change the automatic refresh mode, in the **Mode** field, choose either **Incremental** or **Full**.
- To change the automatic refresh interval, in the **Refresh Interval** field, specify a numeric interval and units of time in minutes, seconds or milliseconds.
- To change the automatic refresh state, in the **Refresh State** field, choose **Paused**, **On** or **Off**.

Click **Apply**.

For more information about automatic refresh of cache groups, see "AUTOREFRESH cache group attribute" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*.

7.10 Specify and Display LOB Data Types

If you are using TimesTen release 11.2.2.2.0 or greater, then you can specify and display LOB data types in your table definitions.

If you want to create a table with a LOB data type, then in the **Create Table** dialog, expand the **Type** column header. You see CLOB, NCLOB, and BLOB as valid data types.

To view the data type for your columns, select the table and then choose the **Columns** tab. You see the LOB data type for your column.

If you want to add a column and specify a LOB data type, in the **Add Column** dialog, expand **Data Type**. You see CLOB, NCLOB, and BLOB among the possible choices.

For PL/SQL objects, you can specify LOB data types as `IN`, `OUT`, and `IN OUT` parameters. In the PL/SQL edit dialogs, for parameter types, choose CLOB, NCLOB, or BLOB.

7.11 Specify the INLINE Attribute for Columns

You can specify the `INLINE` attribute for columns of type `VARCHAR2`, `NVARCHAR2`, and `VARBINARY`.

In the **Create Table** dialog, locate the column header named **Inline**. Click in the checkbox to define the column with the `INLINE` attribute.

You can also specify the `INLINE` attribute when adding a column to a table. Right-click on the table, select **Column**, then select **Add**. The **Add Column** dialog appears. For columns of type `VARCHAR2`, `NVARCHAR2`, and `VARBINARY`, click in the **Inline** checkbox to add the column with the `INLINE` attribute.

7.12 Compute Table Size

If you are using TimesTen Release 11.2.2.0 or greater, you can compute table size information.

To view table size information, you must first compute the table size:

1. Choose **+** to the left of the **Tables** node to view the list of tables.
2. Right-click the name of the table to compute table size information.
3. Select **Table**, then select **Compute Size**.

The **Compute Size** dialog appears. The owner and name of the table are displayed.

4. Click in the checkbox if you want to compute out-of-line sizes.
5. Choose **Apply**.

A **Confirmation** dialog appears.

6. Choose **OK**.

The table size statistics for the table are computed.

To view the SQL for computing the table size, choose the **SQL** tab in the **Compute Size** dialog. You see that a TimesTen built-in procedure called `ttComputeTabSizes` is executed. After this built-in is executed, you can review the table size information for your table. Note that this table size information is on a per table basis.

To compute table sizes for all tables in your database including materialized views, system tables, and tables that are part of cache groups, use the **SQL Worksheet** and execute the command: `Call ttComputeTabSizes (NULL,0);` or to include out-of-line data: `Call ttComputeTabSizes (NULL,1);`

After you compute the table size for one or more tables, you can view the table size information:

1. Choose **+** to the left of the **Tables** node to view the list of tables.
2. Choose the table you want to review table size information.
3. Choose the **Sizes** tab located among the tabs that display the attributes of the table.

Table size information is displayed. Specifically, values for `INLINE BYTES`, `OUT OF LINE BYTES`, and `METADATA BYTES` are shown.

You can choose the **Actions** menu to generate or regenerate the table size. If you choose the **Actions** menu, then **Table**, then select **Compute Size**. The **Compute Size** dialog appears allowing you to compute the table size information.

Choose **Refresh** to refresh the displayed table size data.

You can view table size information for all tables that you have computed table sizes. To view such information, select **TimesTen Reports**, then **Table**, then select **Table Sizes**.

For more information on the `ttComputeTabSizes` built-in procedure, see "ttComputeTabSizes" in *Oracle TimesTen In-Memory Database Reference*.

7.13 Use a REF CURSOR as an OUTPUT Parameter

You can test a PL/SQL function, procedure, or package by defining a REF CURSOR as an OUTPUT parameter in your PL/SQL function, procedure, or package. After you define a REF CURSOR, compile and run your PL/SQL function, procedure, or package. The **Run** dialog appears. Choose **OK**.

After you choose **OK** in the **Run PL/SQL** dialog, you can view the details of the run in the two tabs:

1. **IdeConnections:** Displays the status of your run including any errors encountered.
2. **Output Variables:** Shows the output from the execution of the function, procedure, or package.

7.14 Create an Index on a TimesTen Table

You can create a range index or a bitmap index on a regular table or on a cache table.

To create an index, click the + to the left of the **Tables** node to view the list of tables. Then right-click the name of the table to create an index on and select **Index**, then select **Create Index**.

You can also right-click the **Indexes** node and select **New Index**.

In the **Properties** tab of the **Create Index** dialog, locate the **Table** drop-down menu. From this menu, select the table for the index.

For **Type**:

- To create a non-unique range index, choose **Non-unique**
- To create a unique range index, choose **Unique**
- To create a bitmap index, choose **Bitmap**

For information about the different types of indexes supported in TimesTen, see "CREATE INDEX" in the *Oracle TimesTen In-Memory Database SQL Reference*.

In the **Index** section, specify the columns for the index. The columns can be selected from the **Column Name or Expression** drop-down select list. For each indexed column, from the **Order** select list, you can specify whether the column is to be sorted in ascending or descending order. The default sort order is ascending.

To create a composite index, click + to add columns to the index definition. Click X to remove columns from the index definition.

In the **DDL** tab of the **Create Index** dialog, you can view the CREATE INDEX statement used to create the index.

Click **OK** to create the index.

7.15 Create a Primary Key Constraint

You can create a primary key constraint on a table that does not have a primary key defined. Follow these steps:

1. Click the **+** to the left of the **Tables** node to view the list of tables. Then right-click the name of the table to create a primary key constraint.

2. Select **Constraint**, then select **Add Primary Key**.

The **Add primary key** dialog appears.

3. In the **key index type** field choose either **Use range** to create a primary key using a range index or **Use hash** to create a primary key using a hash index.

If you choose **Use hash**, you see the text **with number of** followed by a drop-down list of either **row** or **pages**. Select either **row** or **pages** and in the box after the **colon (:)**, enter the number of rows or pages. The default selection is **row**.

For more information on hash index sizing, see "Column Definition" in the *Oracle TimesTen In-Memory Database SQL Reference*.

4. In the **Constraint Name** field, enter the constraint name. The name cannot exceed 30 characters in length.

5. In the **Column** field(s), you see a drop-down list of column names. Choose the column name(s) to use for the primary key.

6. Click **Apply**.

The **Confirm** dialog box appears indicating that the primary key constraint has been added.

7. Click **OK**.

The primary key constraint has been added successfully.

To view the SQL statement used to create the primary key constraint, choose the **SQL** tab of the **Add primary key** dialog box.

For information about creating a primary key constraint, see "CREATE TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference* or "ALTER TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference*.

7.16 Change the Attributes of Primary Key

You can change the attributes of a primary key constraint after you have defined a primary key on a table. Follow these steps:

1. Click the **+** to the left of the **Tables** node to view the list of tables. Then right-click the name of the table on which to change the primary key attributes.

2. Select **Constraint**, then select **Change Primary Key Parameters**.

The **Change parameters of a primary key** dialog appears.

3. In the **key index type** field choose either **Use range** to change the primary key constraint from using a hash index to using a range index or **Use hash** to change the primary key constraint from using a range index to using a hash index.

If you choose **Use hash**, you see the text **with number of** followed by a drop-down list of either **row** or **pages**. Select either **row** or **pages** and in the box after the **colon (:)**, enter the number of rows or pages. The default selection is **row**.

For more information on hash index sizing, see "Column Definition" in *Oracle TimesTen In-Memory Database SQL Reference*.

4. Click **Apply**.

The **Confirm** dialog appears indicating that the index(es) for the primary key constraint has been changed.

5. Click **OK**.

The index(es) for the primary key constraint has been changed successfully.

To view the SQL statement used to change the attributes of the primary key constraint, choose the **SQL** tab of the **Change parameters of a primary key** dialog.

For information about changing the attributes of a primary key constraint, see "ALTER TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference*.

7.17 Create a Foreign Key Constraint

You can create a foreign key constraint on a table. Make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Click the **+** to the left of the **Tables** node to expand the list of tables.

2. Right-click the name of the table and select **Constraint**, then select **Add Foreign Key**.

The **Add Foreign Key** dialog displays. The Owner and Name fields are auto-filled and you cannot edit these fields. Prepare to enter the Constraint Name.

3. To enable the **ON DELETE CASCADE** referential action, select the **Cascade delete foreign key** checkbox. If specified, when rows containing referenced key values are deleted from a parent table, rows in child tables with dependent foreign key values are also deleted.

4. In the **Constraint Name** field, enter the name of the foreign key.

5. Select the column on which to apply the foreign key constraint from the **Column Name** drop-down list.

6. Select the table name that the foreign key references from the **Referenced Table Name** drop-down list.

7. Select the column that the foreign key references from the **Referenced Column** drop-down list.

To view the SQL statement that will be used to create the foreign key constraint, choose the **SQL** tab of the **Add Foreign Key** dialog box.

8. Click **Apply**.

The **Confirmation** dialog box displays indicating that the foreign key constraint has been added.

9. Click **OK**.

The foreign key constraint has been added successfully.

For information about creating a foreign key constraint, see "CREATE TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference* or "ALTER TABLE" in the *Oracle TimesTen In-Memory Database SQL Reference*.

7.18 Drop a Constraint

You can drop a constraint from a TimesTen table. Make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Click the **+** to the left of the **Tables** node to expand the list of tables.
2. Right-click the name of the table and select **Constraint**, then select **Drop**.
The **Drop** dialog displays. The Owner and Name fields are auto-filled and you cannot edit these fields. Prepare to select the Constraint.
3. Select the constraint that you want to drop from the **Constraint** drop-down list.
To view the SQL statement that will be used to drop the constraint, choose the **SQL** tab of the **Drop** dialog box.
4. Click **Apply**.
The **Confirmation** dialog box displays indicating that the constrain has been dropped.
5. Click **OK**.
The constraint has been successfully dropped.

7.19 Add an Aging Policy to a Table

To add an aging policy to a regular table or a cache table, under the **Tables** node, right-click the name of the table to add an aging policy to and select **Aging**, then select **Add Usage-based** to add an LRU aging policy. To add a time-based aging policy select **Aging**, then select **Add Time-based**. An aging policy can be added to a cache table only if it is the root table of a cache group.

To add an LRU aging policy to the table, in the **Prompts** tab of the **Add Usage-based** dialog box, specify an aging state by choosing either **On** or **Off** in the **Usage-based aging** field. The default aging state is on. Click **Apply** to add the LRU aging policy to the table.

The TimesTen database memory usage thresholds determine when data starts and stops being deleted from the table. The default memory usage threshold that determines when data starts being deleted from the table is 90%. The default memory usage threshold that determines when data stops being deleted from the table is 80%. The default LRU aging cycle is 1 minute. For information about how to change these thresholds and the LRU aging cycle., see [Change the Memory Usage Thresholds and LRU Aging Cycle](#) (page 7-21).

An LRU aging policy can be added to any regular table, and only to cache tables in an AWT, SWT or user managed cache group that does not have automatic refresh defined.

To add a time-based aging policy to the table, in the **Prompts** tab of the **Add time-based** dialog box, select the column name from **Column** to store the timestamp value indicating when each row was added or most recently updated.

In the **Life-time** field, indicate the length of time in which data that has not been updated is to be kept in the table by specifying a numeric value followed by a unit of minutes, hours or days in the drop-down menu.

In the **Cycle** field, indicate the frequency at which data is to be aged out of the table by specifying a numeric value followed by a unit of minutes, hours or days. The default time-based aging cycle is 5 minutes.

Specify an aging state by choosing either **On** or **Off** in the **Time-based aging** field. The default aging state is on. Then click **Apply** to add the time-based aging policy to the table.

A time-based aging policy can only be added to a table that contains a non-nullable DATE or TIMESTAMP column.

7.20 Drop an Aging Policy from a Table

To drop an existing aging policy from a regular table or a cache table, under the **Tables** node, right-click the name of the table to drop an aging policy from and then select **Aging**, then select **Drop**.

Click **Apply** to drop the aging policy from the table.

7.21 Change the Aging State for the Aging Policy of a Table

To change the aging state of a regular table's or cache table's existing aging policy, under the **Tables** node, right-click the name of the table to change the aging state of and then select **Aging**, then select **Change State On/Off**.

In the **Prompts** tab of the **Change state on/off** dialog, change the aging state by selecting either **On** or **Off** in the **Change aging state** field. Click **Apply** to change the aging state of the table.

7.22 Change the Memory Usage Thresholds and LRU Aging Cycle

To change the memory usage thresholds and the LRU aging cycle, right-click the node of the connection name for the TimesTen database and choose **Change Usage-based Aging Attributes**.

In the **Prompts** tab of the **Changing the usage-based thresholds** dialog box, specify the threshold that determines when data starts being deleted from the tables in the **High usage threshold** field. Specify the threshold that determines when data stops being deleted from the tables in the **Low usage threshold** field. Specify the LRU aging cycle in the **Update Frequency** field. Then click **Apply** to change the memory usage thresholds and the LRU aging cycle.

The new settings apply to all tables that have an LRU aging policy defined.

7.23 Change the Lifetime and Cycle for the Time-Based Aging Policy of a Table

To change the lifetime and cycle of a regular table's or cache table's existing time-based aging policy, under the **Tables** node, right-click the name of the table to change the lifetime and cycle of and select **Aging**, then select **Change Lifetime and Cycle**.

In the **Life-time** field within the **Prompts** tab of the **Change Lifetime and Cycle** dialog box, change the length of time in which data that has not been updated is to be kept in the table by specifying a numeric value followed by a unit of minutes, hours or days.

In the **Cycle** field, change the frequency at which data is to be aged out of the table by specifying a numeric value followed by a unit of minutes, hours or days.

Then click **Apply** to change the lifetime and cycle for the table.

7.24 Grant Object Privileges to a TimesTen Database Element

You can grant privileges to and revoke privileges from the following database objects:

- Regular and cache tables
- Regular and materialized views
- Sequences
- PL/SQL packages, procedures and functions

To grant privileges on an object, right-click the name of the object and select **Privileges**, then select **Grant**.

In the **Users** drop-down menu of the **Grant** dialog box, select the user to grant object privileges to the object. In the **Privileges** section, select the **All** check box or click >> to grant all available object privileges on the object to the selected user. Otherwise, select the individual privileges from the **Available Privileges** list for the privileges to grant to the selected user and then click > to move those privileges into the **Selected Privileges** list. To select multiple privileges, press and hold the CTRL key, and click the desired privileges. Click **Apply** to grant the selected object privileges on the object to the selected user.

For a particular object type, only the available object privileges are shown. For example, DELETE, INDEX, INSERT, REFERENCES, SELECT and UPDATE privileges can be granted on a table. However, only INDEX, REFERENCES and SELECT privileges can be granted on a materialized view.

Only users with the ADMIN system privilege or the owner of an object can grant object privileges on the object.

For more information about the set of privileges available to each type of object, see "Object privileges" in the *Oracle TimesTen In-Memory Database SQL Reference*.

7.25 Revoke Object Privileges from a TimesTen Database Element

To revoke privileges from an object, right-click the name of the object and select **Privileges**, then select **Revoke**.

In the **Users** drop-down menu of the **Revoke** dialog box, select the user to revoke object privileges from the object. In the **Privileges** section, select the **All** check box or click >> to revoke all granted object privileges on the object from the selected user. Otherwise, select the individual privileges from the **Available Privileges** list for the privileges to revoke from the selected user and then click > to move those privileges into the **Selected Privileges** list. To select multiple privileges, press and hold the CTRL key, and click the desired privileges. Click **Apply** to revoke the selected object privileges on the object from the selected user.

The **Users** drop-down menu shows only users who have been granted object privileges on the object. For a particular user, only the object privileges that they have been granted are shown.

Only users with the ADMIN system privilege or the owner of an object can revoke object privileges from the object.

7.26 Load Data into a Table

You can load data into your table using **Import Data**. Before you can load your data, you must export the data into a file. To export the data, right-click on the name of the table and choose **Export**. After successfully exporting your data, you can import the data.

1. Choose **+** to the left of the **Tables** node to view the list of tables.
2. Right-click on the table name.
3. Select **Table**, then select **Import Data**.

An Open dialog appears. Locate the data file to import. This file is the file you previously exported.

4. Choose **Open**.

The data is loaded into the table.

7.27 Start and Stop the TimesTen Replication Agent

The replication agent process transmits committed updates on replicated tables in a master TimesTen database to the corresponding replicated tables in a subscriber TimesTen database. The replication agent is also responsible for asynchronously propagating committed updates on cache tables in AWT cache groups to the corresponding Oracle tables. After you create a replication scheme or an AWT cache group, you must start the replication agent if it is not running.

To start the replication agent, right-click the node of the connection name for the TimesTen database and select **Start/Stop Replication Agent**.

The **Starting/stopping the replication agent** dialog appears. In the **Prompts** tab, select **Start agent**. Click **Apply**.

If there are no replication schemes or AWT cache groups in your TimesTen connection, the replication agent is not started.

To stop a running replication agent process, right-click the node of the connection name for the TimesTen database and select **Start/Stop Replication Agent**.

The **Starting/stopping the replication agent** dialog appears. In the **Prompts** tab, select **Stop agent**. Click **Apply**.

See Also:

"Starting and stopping the replication agents" in the *Oracle TimesTen In-Memory Database Replication Guide* and "Managing the replication agent" in the *Oracle TimesTen Application-Tier Database Cache User's Guide*

7.28 Update the Table and Column Statistics Used By the TimesTen Query Optimizer

The query optimizer uses statistics stored in TimesTen database system tables to determine the optimal execution plan for a statement. TimesTen stores table level statistics in the `SYS.TBL_STATS` table. Column level statistics, such as the minimum

and maximum value and the number of unique values in a column, are stored in the `SYS.COL_STATS` table.

TimesTen does not compute table and column statistics as updates occur on regular and cache tables. Instead, TimesTen updates statistics when an explicit request is made.

You can update statistics for all tables owned by a user or for a specific table owned by an user.

To update statistics for all tables, right-click the **Tables** node and select **Statistics**, then select **Update**.

To update statistics for a specific table, click the **+** to the left of the **Tables** node. Right-click the name of the table and select **Statistics**, then select **Update**.

In the **Prompts** tab of the **Update statistics** dialog, choose **Invalidate referenced commands** to invalidate the execution plans of statements that reference the table for which statistics are updated. When you invalidate the execution plan of a statement, TimesTen recompiles or reprepares that statement upon its next execution.

For **Interval type for table statistics**, choose **Complete interval** to divide the rows of the table into two or more intervals and compute statistics on each interval, or **Single interval** to compute statistics on the entire set of rows as a single interval. You must define a range index on the table to compute complete interval statistics.

Click **Apply**.

To view the optimizer statistics of a table, click the name of the table.

In the **Statistics** tab of the table pane, the top section shows the table statistics such as:

- The number of rows in the table
- A timestamp indicating when statistics were most recently updated for the table

The bottom section shows the column statistics. For each column, the following information is displayed:

- The name of the column
- The number of intervals the data is divided into to compute statistics
- Total number of NULL values
- Total number of non-NULL unique values
- Total number of rows in each interval

For each interval of each column, SQL Developer displays the following:

- The number of unique values other than the most frequently occurring value
- The number of rows that contain a value other than the most frequently occurring value
- The number of rows that contain the most frequently occurring value
- The minimum value
- The maximum value
- The most frequently occurring value

For more information about query optimizer statistics, see "Statistics" in the *Oracle In-Memory Database Cache Introduction*, "When optimization occurs" in the *Oracle TimesTen In-Memory Database Operations Guide* or "ttOptUpdateStats" in the *Oracle TimesTen In-Memory Database Reference*.

7.29 Export the Table Statistics Used By the TimesTen Query Optimizer

You can generate a set of SQL statements that allow you to restore the table statistics to the current state. You can generate SQL statements for a single table or for user tables that the connection user has permission to access. After you export your table statistics to a .sql file, you can import your table statistics by opening the .sql file in SQL Developer and running the script file.

To export the current TimesTen table statistics, make sure you are on the main SQL Developer page.

1. To save table statistics for a single table, right-click the name of the table and select **Statistics**, then select **Export**. To save table statistics for all user tables, right-click **Tables** and select **Statistics**, then select **Export**.

The **Export Optimizer Statistics** dialog displays. Locate the **Save** button at the bottom right of the dialog.

2. Click **Save**.

The **Save Optimizer Statistics Export Commands** dialog displays.

3. Select the directory where you want to save the current TimesTen table statistics.
4. In the **File Name** field, define the name of your current TimesTen table statistics.

The default filename of the current TimesTen table statistics is *connection_name-exportstatistics-YYYYMMDDMISS.sql*, where *connection_name* is the name of the connection. *YYYYMMDDHHMISS* is a timestamp of when the recommendations were made, where *YYYY* is the year, *MM* the month, *DD* the day, *HH* the hour, *MI* the minutes, and *SS* the seconds.

5. Click **Save**.

TimesTen saves your current TimesTen table statistics.

7.30 Create a TimesTen Performance Metrics Snapshot

A snapshot is a collection of performance metrics from a database. You can create a report comparing two snapshots to analyze the performance of your database.

To create a new snapshot of a database, make sure you are on the main SQL Developer page and that you have expanded the TimesTen connection in the connections navigator.

1. Right-click the name of the database and select **TimesTen Performance Metrics**, then select **Create Snapshot**.

The **New TimesTen Metrics Snapshot** dialog displays.

2. In the **Prompts** tab of the **New TimesTen Metrics Snapshot** dialog, select the desired capture level from the **Capture Level** drop-down list:
 - **TYPICAL**: For a typical set of metrics. This level is appropriate for most purposes. This is the default capture level.

- **NONE:** For metrics outside of SYS . SYSTEMSTATS only.
- **BASIC:** For a minimal basic set of metrics.
- **ALL:** For all available metrics.

Make sure to use the same capture level for any two snapshots to be used in a report. An error dialog displays if you attempt to generate a report that compares two snapshots with different capture levels.

You are now ready to create a new snapshot.

3. Click **Apply** to create the snapshot.

A confirmation dialog shows the ID number of the snapshot. TimesTen generates the snapshot ID.

4. Click **OK**.

TimesTen creates the snapshot with the specified snapshot ID. The snapshot ID is useful for creating TimesTen metrics reports.

7.31 Drop a TimesTen Performance Metrics Snapshot

Consider dropping unneeded snapshots to free up permanent space. If you need to save more snapshots of your database, consider editing the TimesTen performance metrics configuration. See "Editing the TimesTen performance metrics configuration" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

To drop a new snapshot from a database, make sure you are on the main SQL Developer page.

1. Right-click the name of the database and select **TimesTen Performance Metrics**, then select **Drop Snapshot**.
2. By default, the **Prompts** tab of the **Drop TimesTen Metrics Snapshot(s)** dialog is selected. In the **Selector** column, select the snapshots you want to drop.

You can also click on the column header to sort the table based on the column. A description of each column follows:

- **Selector**

The check box you select to determine the snapshots that you want to drop.

- **Snapshot**

The snapshot ID of the snapshot.

- **Timestamp**

The timestamp of when the snapshot was captured.

- **Capture Level**

The capture level of the snapshot. This value can be TYPICAL, NONE, BASIC, or ALL.

To select all snapshots, click the **Select all** check box.

3. Once you have selected the snapshots that you want to drop, click the **Apply** button.

A Confirm dialog displays asking you if you want to proceed with dropping the selected snapshots.

4. Click **Yes**.

A confirmation dialog displays a message indicating that one or more TimesTen Metrics snapshots were dropped.

5. Click **OK**.

You successfully dropped the TimesTen metrics snapshots.

7.32 Edit the TimesTen Performance Metrics Configuration

TimesTen performance metrics configuration allows you to modify the maximum snapshot count and the maximum snapshot retention size. These configuration settings are useful both in preventing your database from running out of permanent space and in allowing you to save more snapshots.

Once the maximum snapshot count or the maximum snapshot retention size is reached, TimesTen removes the oldest snapshots to free up permanent space. By default, the maximum snapshot count is 256 and the maximum snapshot retention size is 256 MB.

To edit the TimesTen performance metrics configuration for a database, make sure you are on the main SQL Developer page.

1. Right-click the name of the database and select **TimesTen Performance Metrics**, then select **Edit Configuration**.

The **TimesTen Metrics Configuration** dialog displays and you see the current configuration for the TimesTen metrics.

2. The **Prompts** tab of the **TimesTen Metrics Configuration** dialog is selected. In the **Maximum Snapshot Count** field, specify a numeric value to indicate the maximum number of snapshots. By default, the maximum number of snapshots is 256.
3. In the **Maximum Snapshot Retention Size (MB)** field, specify the maximum total size of all stored snapshots, in megabytes. By default, the maximum snapshot retention size is 256 MB.
4. Click **Apply**.

You have applied the new TimesTen Metrics configuration settings.

7.33 Create a TimesTen Performance Metrics Report

TimesTen performance metrics reports use the data from two specified snapshots of TimesTen metrics to produce a report in HTML format. TimesTen metrics reports include information for each metric. For more information about the various metrics displayed in the report, see "Report examples" in the *Oracle TimesTen In-Memory Database Reference*.

To generate a new report for a database, make sure you are on the main SQL Developer page.

1. Right-click the name of the database and select **TimesTen Performance Metrics**, then select **Report**.

The **Create TimesTen Metrics Report** dialog displays.

2. By default, the **Report generator** tab of the **Create TimesTen Metrics Report** dialog is selected. In the **Selector** column, select two snapshots to compare.

You can also click on the column header to sort the table based on the column. A description of each column follows:

- **Selector**

The check box that determines the snapshots that are compared in the metrics report. Make sure to only select two snapshots with the same capture level. An error dialog displays if you attempt to generate a report that compares more than two snapshots.

- **Snapshot**

The snapshot ID of the snapshot.

- **Timestamp**

The timestamp of when the snapshot was captured.

- **Capture Level**

The capture level of the snapshot. This value can be TYPICAL, NONE, BASIC, or ALL.

3. Once you have selected two snapshots with the same capture level, click the **Generate report** button. An error dialog displays if you attempt to generate a report that compares snapshots with different capture levels.

TimesTen generates your TimesTen metrics report. You can view the generated TimesTen metrics report in the pane located below the **Generate report** button.

If you do not want to save the report, click **Close**. If you want to save the report, see [Save a report](#) (page 7-28).

7.33.1 Save a report

If you want to save the report for future reference, follow these steps:

1. Click **Save**.

The Save dialog displays.

2. Select the directory where you want to save the TimesTen metrics report.

3. In the **File Name** field, define the name of your TimesTen metrics report.

The default filename of the TimesTen metrics report is `ttstats-snapshot_id_1-snapshot_id_2-YYYYMMDDHHMISS.html`, where `snapshot_id_1` and `snapshot_id_2` are the snapshot IDs of the two snapshots you are comparing in your report. `YYYYMMDDHHMISS` is a timestamp of when the report was generated, where `YYYY` is the year, `MM` the month, `DD` the day, `HH` the hour, `MI` the minutes, and `SS` the seconds.

4. Click **Save**.

SQL Developer saves your report.

7.34 Capture Data for the TimesTen Index Advisor at the Connection Level

Before attempting to capture data for the TimesTen Index Advisor, make sure your table statistics have been updated in the past 24 hours. Up-to-date table statistics provide the most up-to-date statistics for data collection and allow statements to be prepared with the optimized query plan. For more information about updating table statistics, see "Updating the table and column statistics used by the query optimizer" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

You have the following options:

- [Capture data from a SQL worksheet](#) (page 7-29)
- [Capture data from a SQL workload script](#) (page 7-31)

7.34.1 Capture data from a SQL worksheet

You can capture data for the TimesTen Index Advisor at the connection level from the SQL Developer SQL worksheet. Make sure your SQL workload is in the SQL Developer SQL worksheet.

1. Click the **TimesTen Index Advisor** button from the SQL worksheet menu bar.
 - If your table statistics are outdated, an information dialog displays information about the outdated tables. TimesTen recommends that you update your table statistics.
 - If your table statistics are up to date, the **Index Advisor Configuration** dialog displays.
2. To update your table statistics, click **Cancel** and follow the steps in "Updating the table and column statistics used by the query optimizer" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*. To proceed without updating your table statistics, click **OK**.

The **Index Advisor Configuration** dialog displays.

3. Select the desired capture mode from the **Capture mode** drop-down list:
 - **Prepare SQL:** TimesTen prepares but does not execute the SQL commands from the SQL worksheet. The TimesTen Index Advisor makes recommendations using computed statistics and query plan analysis. This is the default capture mode.
 - **Execute SQL:** TimesTen executes the SQL commands from the SQL worksheet. The TimesTen Index Advisor makes recommendations using the actual execution of the SQL commands. This capture mode may take longer than the Prepare SQL mode because TimesTen has to complete SQL execution.

You are ready to choose if you want to use optimizer hints.

4. The **Include optimizer hints for Oracle BI server** checkbox allows you to use optimizer hints that direct the TimesTen query optimizer to generate a specific execution plan. These optimizer hints are recommended queries generated by the Oracle BI server. If you do not want to use optimizer hints, click **Start**.

The **Index Advisor Configuration** dialog closes. A **TimesTen index advisor** pane displays at the bottom of the SQL worksheet.

5. In the **Selector** column, select the index recommendations you want to create.
You can also click on the column header to sort the table based on the column. A description of each column follows:
 - **Selector**
The check box you select to determine the index recommendations that you want to create.
 - **Index Recommendation**
The `CREATE INDEX` statement that the TimesTen index advisor recommends.
 - **Affected Statement Count**
The number of statements that benefit from the recommendation.
 - **Created**
Shows if you have already created the index recommendation. This value can be Yes or No.

If you want to select all index recommendations, click the **Select all** check box.

If you want to review the SQL commands that are being evaluated by the TimesTen index advisor, click the **SQL** tab. The **SQL** tab is located at the top right of the **TimesTen index advisor** pane.

6. Once you select the index recommendations that you want to create, click the **Create Selected Indexes** button.
The **Create Selected Indexes** dialog displays. Locate the **Details >>** button.
7. Click **Details >>**.
A details pane expands that shows progress information about creating the selected indexes. Locate the **Close** button at the bottom of the dialog.
8. Click **Close**.

The **Create Selected Indexes** dialog closes. TimesTen creates the indexes. If you want to save the index recommendations, see [Saving index recommendations from a SQL worksheet](#) (page 7-30).

See Also:

"Including optimizer hints for Oracle BI server in a SQL worksheet" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*

7.34.1.1 Saving index recommendations from a SQL worksheet

To save index recommendations for future reference, follow these steps:

1. Click **Save Indexes**. TimesTen saves all index recommendations. TimesTen does not save individual index recommendations.
The **Save Indexes** dialog displays.

2. Select the directory where you want to save the TimesTen index advisor recommendations.
3. In the **File Name** field, define the file name of your TimesTen index advisor recommendations.

The default filename of the TimesTen index advisor recommendations is `connection_name-indexadvice-YYYYMMDDMMISS.sql`, where `connection_name` is the name of the connection. `YYYYMMDDHHMISS` is a timestamp of when the recommendations were made, where `YYYY` is the year, `MM` the month, `DD` the day, `HH` the hour, `MI` the minutes, and `SS` the seconds.

4. Click **Save**.

TimesTen saves the index recommendations.

The saved file header contains the following information:

- The timestamp of when TimesTen generated the report.
- The version of SQL Developer.
- The version of the TimesTen database.
- The name of the TimesTen connection.

For example:

```
-- This file was generated by SQL Developer at 2013-04-29 11:08:45
-- SQL Developer version 4.0.0.11.51
-- Database version: 11.02.02.0005 Oracle TimesTen IMDB version 11.2.2.5.0
-- Connection name: sampledb_1122
```

7.34.2 Capture data from a SQL workload script

You can capture data for the TimesTen Index Advisor at the connection level using the SQL queries from a SQL workload script file. Make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Right-click the name of the database and select **Index Advisor**, then select **Connection Level Capture**.
 - If your table statistics are outdated, an information dialog displays information about the outdated tables. TimesTen recommends that you update your table statistics.
 - If your table statistics are up to date, the **Index Advisor Configuration** dialog displays.
2. To update your table statistics, click **Cancel** and follow the steps in "Updating the table and column statistics used by the query optimizer" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*. To proceed without updating your table statistics, click **OK**.

The **Index Advisor - Connection Level Capture** dialog displays.

3. In the **Configuration** tab of the **Index Advisor - Connect Level Capture** dialog, select the desired capture mode from the **Capture mode** drop-down list:
 - **Prepare SQL**: TimesTen prepares but does not execute the SQL commands from the workload script. The TimesTen Index Advisor makes

recommendations using computed statistics and query plan analysis. This is the default capture mode.

- **Execute SQL:** TimesTen executes SQL commands from the workload script. The TimesTen Index Advisor makes recommendations using the actual execution of the workload script. This capture mode may take longer than the Prepare SQL mode because TimesTen has to complete SQL execution.

You are ready to specify your SQL workload script. Locate the **Browse** button that is to the right of the **Workload script** field.

4. Click **Browse**.

The **Open** dialog displays.

5. Select the directory where you saved your SQL workload script.
6. Select the SQL workload script for which you want to generate index recommendations.
7. Click **Open**.

The TimesTen index advisor is now ready to collect data for your SQL workload. Locate the **Start** button.

To review the SQL commands that are being evaluated by the TimesTen index advisor, click the **SQL** tab. The **SQL** tab is located at the top of the **Index Advisor - Connection Level Capture** dialog.

8. The **Include optimizer hints for Oracle BI server** checkbox allows you to use optimizer hints that direct the TimesTen query optimizer to generate a specific execution plan. These optimizer hints are recommended queries generated by the Oracle BI server. If you do not want to use optimizer hints, click **Start**.

The **Index advisor assistant progress** dialog displays. Locate the **Details >>** button.

If the TimesTen index advisor does not have any recommendations, the **Index Recommendation Feedback** dialog displays. Close this dialog to return to the main SQL developer page.

9. Click **Details >>**.

A details pane expands that shows progress information about the index advisor connection level capture. Locate the **Close** button at the bottom of the dialog.

10. Click **Close**.

The **Index advisor assistant progress** dialog closes. Locate the **Index Recommendations** tab.

11. Click the **Index Recommendations** tab.

The **Index Recommendations** tab of the **Index Advisor - Connection Level Capture** dialog displays.

12. In the **Selector** column, select the index recommendations that you want to create.

You can also click on the column header to sort the table based on the column. A description of each column follows:

- **Selector**

The check box you select to determine the index recommendations that you want to create.

- **Index Recommendation**

The `CREATE INDEX` statement that the TimesTen index advisor recommends.

- **Affected Statement Count**

The number of statements that benefit from the recommendation.

- **Created**

Shows if you have already created the index recommendation. This value can be Yes or No.

If you want to select all index recommendations, click the **Select all** check box.

13. Once you have selected the index recommendations that you want to apply, click the **Create Selected Indexes** button.

The **Creating Selected Indexes** dialog displays. Locate the **Details >>** button.

14. Click **Details >>**.

A details pane expands that shows progress information about creating the selected indexes. Locate the **Close** button at the bottom of the dialog.

15. Click **Close**.

The **Creating Selected Indexes** dialog closes. TimesTen created the indexes. If you want to save the index recommendations, see "Saving index recommendations" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

See Also:

For more information on using optimizer hints, see "Including optimizer hints for Oracle BI server for a SQL workload script" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

7.35 Capture Data for the TimesTen Index Advisor at the Database Level

Before attempting to capture data for the TimesTen Index Advisor, make sure your table statistics have been updated in the past 24 hours. Up-to-date table statistics provide the most up-to-date statistics for data collection and allow statements to be prepared with the optimized query plan. For more information about updating table statistics, see "Updating the table and column statistics used by the query optimizer" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

To collect data for the TimesTen index advisor at the database level, make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Right-click the name of the database and select **Index Advisor**, then select **Database Level Capture**.

- If your table statistics are outdated, an information dialog displays information about the outdated tables. It is recommended that your table statistics are up to date.

- If your table statistics are up to date, the **Index Advisor - Database Level Capture** dialog displays.
2. To update your table statistics, click **Cancel** and follow the steps in "Updating the table and column statistics used by the query optimizer" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*. If you want to proceed without updating your table statistics, click **OK**.
The **Index Advisor - Database Capture** dialog displays. Locate the **Start** button at the bottom of the dialog.
 3. In the **Control** tab of the **Index Advisor - Connect Level Capture** dialog, click **Start**.
The **Index Advisor Database Capture In Progress** dialog displays. Locate the **Details >>** button.
 4. Click **Details >>**.
A details pane expands that shows progress information about the index advisor database level capture. Locate the **Close** button at the bottom of the dialog.
 5. Click **Close**.
The **Index Advisor Database Capture In Progress** dialog closes.
 6. Once you have captured your desired SQL workload, click **Stop**.
If the TimesTen index advisor has recommendations, the **Index Advisor Database Capture In Progress** dialog displays.
If the TimesTen index advisor does not have any recommendations, the **Index Recommendation Feedback** dialog displays. You may close this dialog to return to the main SQL developer page.
 7. If the TimesTen index advisor has recommendations, click **Details >>**.
A details pane expands that shows progress information about the index advisor database capture. Locate the **Close** button at the bottom of the dialog.
 8. Click **Close**.
The **Index Advisor Database Capture In Progress** dialog closes. The **Index Advisor Database Capture** dialog shows information about the completed database capture. Locate the **Index Recommendations** tab at the top of the dialog.
 9. Click the **Index Recommendations** tab.
The **Index Recommendations** tab of the **Index Advisor - Database Capture** dialog displays.
 10. In the **Selector** column, select the index recommendations that you want to create.
You can also click on the column header to sort the table based on the column. A description of each column follows:
 - **Selector**
The check box you select to determine the index recommendations that you want to create.
 - **Index Recommendation**

The `CREATE INDEX` statement that the TimesTen index advisor recommends.

- **Affected Statement Count**

The number of statements that benefit from the recommendation.

- **Created**

Shows if you have already created the index recommendation. This value can be Yes or No.

If you want to select all index recommendations, click the **Select all** check box.

11. Once you have selected the index recommendations that you want to apply, click the **Create Selected Indexes** button.

The **Creating Selected Indexes** dialog displays. Locate the **Details >>** button.

12. Click **Details >>**.

A details pane expands that shows progress information about creating the selected indexes. Locate the **Close** button at the bottom of the dialog.

13. Click **Close**.

The **Creating Selected Indexes** dialog closes. The indexes are created. If you want to save the index recommendations, see "Saving index recommendations" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*.

7.36 View the TimesTen Index Advisor Capture State

The TimesTen index advisor capture state allows you to view the status of connection and database level captures. To view the TimesTen index advisor capture state for a database, make sure you are on the main SQL Developer page.

1. Right-click the name of the database and select **Index Advisor**, then select **Index Advisor Capture State**.

The **Index Advisor Capture State** dialog displays and you see the current status of connection and database level captures.

You can also click on the column header to sort the table based on the column. A description of each column follows:

- **Connection ID**

The connection ID of the connection that initiated the TimesTen index advisor capture.

- **Capture Level**

The capture level of the TimesTen index advisor capture. This value can be Connection or Database.

- **Capture Mode**

The capture mode of the TimesTen index advisor capture. This value can be Execute SQL or Prepare SQL.

- **Capture State**

The capture state of the TimesTen index advisor capture. This value can be In Progress or Completed.

- **Prepared Count**

The number of prepared SQL statements during the capture period.

- **Execute Count**

The number of executed SQL statements during the capture period.

- **Start Time**

The timestamp of when the index advisor capture began. The timestamp format is *YYYY-MM-DD HH:MI:SS.FF*, where *YYYY* is the year, *MM* the month, *DD* the day, *HH* the hour, *MI* the minutes, *SS* the seconds, and *FF* the milliseconds.

- **End Time**

The timestamp of when the index advisor capture ended. The timestamp format is *YYYY-MM-DD HH:MI:SS.FF*, where *YYYY* is the year, *MM* the month, *DD* the day, *HH* the hour, *MI* the minutes, *SS* the seconds, and *FF* the milliseconds. This value is empty if the **Capture State** is in In Progress.

2. Click **Refresh** to update the table with the most current information of the TimesTen index advisor captures.

The **Index Advisor Capture State** dialog is refreshed with the most recent capture state information.

7.37 Load Data from an Oracle Database into a TimesTen Database

SQL Developer allows you to load data from an Oracle database into a TimesTen database without creating a cache grid, cache group, or cache table.

Before attempting to load data from an Oracle database into a TimesTen database:

- Make sure you are using Oracle TimesTen In-Memory Database Release 11.2.2.4 (or later).
- Make sure you have specified the Oracle connection information (Oracle Database Password and Oracle Net Service Name) for your TimesTen SQL developer connection.

To load data from an Oracle database into a TimesTen database, make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Right-click the **Tables** option and select **Load from Oracle Database Tables**.

The **Load from Oracle Database Tables** dialog displays. Locate your Oracle database schema name from the navigation tree on the left of the **Load from Oracle Database Tables** dialog.

2. Click the + located next to your Oracle database schema name to expand the Oracle database tables list.

Your Oracle database schema expands showing a list of available Oracle database tables. Locate the Oracle database tables you want to load into the TimesTen database.

3. Right-click the name of the Oracle database table that you want to load into the TimesTen database and select **Load**. You can repeat this step, including for the same Oracle database table, to load multiple Oracle database tables.

If you want to load all of the Oracle database tables, right-click the Oracle database schema name and select **Load All**. The TimesTen Username dialog displays. Enter the TimesTen user for which you want to create the tables and click **OK**.

Your selected tables are added to the right pane and are displayed in a table format.

4. Review the selected Oracle database tables and edit any necessary fields. You can edit a field by double clicking it. A description of each column follows:
 - **TimesTen Username**
The owner of the TimesTen table that you are creating. By default, this is the TimesTen connection user. You can edit this column. An error message is displayed if the specified table owner does not exist or the TimesTen connection user does not have privileges to access the specified table owner.
 - **TimesTen Table Name**
The name of the TimesTen table that you are creating. By default, this is the same table name as the Oracle database table. You can edit this column.
 - **Query to Run on the Oracle database**
The SQL query executed on the Oracle database to generate the desired result set. By default, this is a `SELECT * FROM oracledb_tbl` query, where `oracledb_tbl` is the Oracle database table from which you are loading data. You can edit this column.
 - **Row Count**
The number of rows that TimesTen loads from the Oracle database table. This is the row count of the result set from the "Query to Run on the Oracle database" field. By default, this value is not counted.

To populate the **Row Count** field for a specific Oracle database table, right-click on the **Row Count** field for your table, then select **Count**.

To populate the **Row Count** fields for all the selected Oracle database tables, right-click on a **Row Count** field of any table, then select **Count All**. A confirmation dialog displays. Click **Yes** to continue with the row count of all of your previously selected Oracle database tables.
 - **Create Table Statement**
The SQL statement that TimesTen uses to create the table to load your Oracle database data. You cannot edit this column. Review the `CREATE TABLE` statement by double clicking on the **Create Table Statement** field.

If a TimesTen table with the same name already exists in the TimesTen database, the data from the Oracle database table is appended to the existing TimesTen table. TimesTen does not load table constraints from the Oracle database table. An error dialog displays if a TimesTen table with the same name but different table structure already exists in the TimesTen database.
 - **Number of Parallel Loads**
The number of parallel loads that will be used to load the Oracle database table into the TimesTen database. If the table contains less than 500,000 rows the default value is one. If the Oracle database table contains more than 500,000 rows the default value is two. You can edit this column.

TimesTen recommends the use of multiple parallel loads when you are loading a large number of rows. Multiple parallel loads speeds up the process of loading rows by having each parallel load import a subset of rows from the Oracle database table into the TimesTen database.

- **Number of Threads per Parallel Load**

The number of parallel threads that will be used to load the Oracle database table into the TimesTen database. By default, the value is four. You can edit this column. An error dialog displays if you attempt to use a number of parallel threads that is less than two.

- **Status**

The status of the data loading operation from the Oracle database to the TimesTen database. This value is empty before starting the load operation. Once you start the load tables from the Oracle database operation, this value can be IN PROGRESS, DONE, or FAILED.

Locate the **Index Creation Policy** at the bottom of the **Load from Oracle Database Tables** dialog.

5. Select the desired index creation policy from the **Index Creation Policy** drop-down list:

- **Copy all indexes from peer Oracle database table:** TimesTen creates all of the indexes that exist on the peer Oracle database table. This is the default index creation policy.
- **Copy unique indexes only from peer Oracle database table:** TimesTen creates all of the unique indexes that exist on the peer Oracle database table.
- **Do not copy indexes from peer Oracle database table:** TimesTen does not create any indexes.

Locate the **Load** button in the bottom right of the **Load from Oracle Database Tables** dialog.

6. Once you edit all of your desired fields, click the **Load** button.

A progress dialog displays. Verify that each field of the **Status** column of the **Load from Oracle Database Tables** dialog is marked as DONE.

If any **Status** field is marked as something other than DONE, make sure that all the data you entered is correct. Any status that is not marked as DONE indicates that an error occurred.

If all of the **Status** fields are marked as DONE, your TimesTen tables are created and populated with the Oracle database table data. Double click the DONE status to see more information regarding the imported indexes.

Locate the **Close** button in the bottom right of the **Load from Oracle Database Tables** dialog.

7. Click **Close**.

You have successfully loaded data from the Oracle database to the TimesTen database.

See Also:

For information on valid Oracle Database SQL queries, see *Oracle Database SQL Language Reference*.

"Defining a TimesTen database named connection" in the *Oracle SQL Developer Oracle TimesTen In-Memory Database Support User's Guide*

7.38 Load Data from an Oracle Database Table into an Existing TimesTen Table

To load data from an Oracle database table into an existing TimesTen table, follow these steps. Make sure you are on the main SQL Developer page and that your TimesTen connection is expanded.

1. Click the + located next to **Tables** to expand the TimesTen tables list.

The TimesTen tables list expands.

2. Right-click the name of the table and select the **Table** option, then select the **Load from Oracle** option.

The **Load Tables from Oracle** dialog displays. The Schema name, TimesTen Username, and TimesTen Table Name are auto-filled and you cannot edit these fields. Prepare to enter the SQL query and the Parallel Thread Count.

3. In the **Select Query** field, enter the SQL query to execute on the Oracle database. The result of the SQL query is loaded into your TimesTen table.
4. In the **Parallel Thread Count** field, enter the number of parallel threads that will be used to load the Oracle database table into the TimesTen database. An error dialog displays if you attempt to use a number of parallel threads that is less than two.
5. Click **Load**.

A progress dialog displays. Once TimesTen loads the data from the Oracle database table, the progress dialog closes. If you execute the query on a small Oracle database table, the progress dialog may close almost immediately.

See Also:

For more information on valid Oracle Database SQL queries, see *Oracle Database SQL Language Reference*.

SQL Developer Accessibility Information

Note:

Unless otherwise indicated, the SQL Developer accessibility information presented here also applies to Data Modeler and Data Miner.

This topic provides information on the accessibility features of Oracle SQL Developer. It includes the following:

- [About Oracle SQL Developer Accessibility](#) (page 8-1)
- [Using a Screen Reader and Java Access Bridge with Oracle SQL Developer](#) (page 8-1)
- [Oracle SQL Developer Features that Support Accessibility](#) (page 8-1)
- [Recommendations for Customizing Oracle SQL Developer](#) (page 8-4)
- [Highly Visual Features of Oracle SQL Developer](#) (page 8-5)

8.1 About Oracle SQL Developer Accessibility

It is our goal to make Oracle Products, Services, and supporting documentation accessible to the disabled community. Oracle SQL Developer supports accessibility features. For additional accessibility information for Oracle products, see the Oracle Accessibility Program page at:

<http://www.oracle.com/accessibility/>

8.2 Using a Screen Reader and Java Access Bridge with Oracle SQL Developer

In order for assistive technologies, like screen readers, to work with Java-based applications and applets, the Windows-based computer must also have Sun's Java Access Bridge installed. See *Oracle SQL Developer Installation Guide* for the screen reader setup procedure, and for the recommended minimum technology stack.

8.3 Oracle SQL Developer Features that Support Accessibility

Oracle SQL Developer provides features that are designed to support accessibility.

8.3.1 Keyboard Access

Oracle SQL Developer features support keyboard access to SQL Developer functionality; a summary is provided in this section. The following menu and toolbar functionality is provided through keyboard access:

- Users can navigate to and invoke all menu items.
- All toolbar functions are accessible through menu items.
- All menus and menu items have unique and functioning mnemonic keys.
- All context menus within the navigators and source editor can be invoked.
- Frequently used menu items have unique accelerator keys.

The following functionality is available in SQL Developer windows:

- Navigate between all open windows, to all nodes within a window or pane, and between tabs in a window.
- Set focus in a window or pane.
- Invoke all controls within a window or pane, and perform basic operations.
- Navigate and update properties in the Property Inspector.
- Use Completion Insight and Code Templates in the source editor.
- Invoke context sensitive help topics, navigate to and open all help topics, and navigate between the navigation and viewer tabs.
- Open, close, dock, undock, minimize, restore, and maximize the applicable SQL Developer window.

Tips:

- You can press Escape to move the focus from the current dockable window to the last active editor. Press Shift+Escape to move the focus and also close the current window.
 - You can press Shift+F10 to open the context menu for any window. Use the Down Arrow and Up arrow keys to select a command and press Enter, or use the accelerators to invoke a command on the context menu.
-
-

The following functionality is available in Oracle SQL Developer dialogs and wizards:

- Users can navigate to and invoke all controls within all wizards and dialogs.
- The order in which the Tab key causes focus to flow is consistent and logical.
- Mnemonic keys are provided for controls where appropriate.

8.3.2 Screen Reader Readability

Here is a summary of screen readability in SQL Developer, when it is used with a screen reader.

When used with menus and toolbars:

- All menus and menu items are read.
- All toolbar items, including the Navigator toolbar items, are read.
- The hint text on all toolbar items is read.

When used with SQL Developer windows:

- All open windows are read.
- All components within each window, including tabs, are read.
- Status text at the bottom of the IDE, and within the source editor, is read.

When used with dialogs and wizards:

- All controls within all wizards and dialogs are read.
- Hint text is read.

8.3.3 Flexibility in Font and Color Choices

The user interface in SQL Developer improves usability for people who are visually impaired by offering flexibility in color and font choices. The following font and color features are included:

- Users can specify both the font and the size in which the font displays for code editors.
- All features of the product have black text on a white or gray background.
- Colored text, underlining, or images are never used as the only method of conveying information.

8.3.4 No Audio-only Feedback

In SQL Developer, there is no situation in which the only feedback a user receives is audible feedback. All audible feedback is accompanied by a visual indicator. For example, a prompt accompanies the bell sound that occurs when an error or illegal action has taken place.

8.3.5 No Dependency on Blinking Cursor and Animation

SQL Developer makes minimal use of a blinking cursor and animation:

- No features in SQL Developer use blinking indicators, with the exception of the cursor in the source editor.
- No features rely on animated sequences.

8.3.6 Screen Magnifier Usability

The SQL Developer user interface works well with screen magnifiers. All features of the product can be magnified by a screen magnifier.

8.3.7 How to Change the Editor or Tabbed View of a File

When you press Enter on a node in the Connections Navigator, you open the default editor for that file. To switch to the different editors and views available for a document; you can use the Alt+Page Up and Alt+Page Down accelerators to invoke the **Navigate > Go to Window > Right Editor** and **Navigate > Go to Window > Left Editor** menu commands, respectively.

8.3.8 How to Read Text in a Multiline Edit Field

To have the text in a multiline edit field read by a screen reader, you can select text by holding down the Shift key while moving the cursor either up or down with the Arrow keys, depending on the initial cursor position.

8.3.9 How to Read the Line Number in the Source Editor

To have the line number read by a screen reader while you are editing a file in the source editor, you can press Ctrl+G.

8.3.10 How to Access Exception Stack HTML Links and Generated Javadoc Links in the Log Window

After generating exception stack HTML links or Javadoc links in the Log window, they will not be recognized as links, but read as plain text by a screen reader. To access the links, set the cursor focus to the Log window. Right-click or press Shift+F1 and select **Save As** from the context menu. Save the contents of the Log window as an HTML file. Add the saved HTML file to a project or application workspace as a resource.

8.4 Recommendations for Customizing Oracle SQL Developer

SQL Developer provides a number of customization features that enable users to specify their requirements for keyboard usage, display attributes of SQL Developer, and timing where appropriate. All customization features are organized within the Preferences dialog. For maximum usability and to accommodate your needs, you should consider changing any of the following from the defaults to a more usable customized setting.

8.4.1 How to Customize the Accelerators Keys

You can add and change the default accelerator keys for Oracle SQL Developer in the **Tools > Preferences > Shortcut Keys** page. You can also load preset keymaps that you are accustomed to using.

8.4.2 How to Pass a Conflicting Accelerator Key to Oracle SQL Developer

In addition to changing the mapped accelerator keys, you can pass a conflicting accelerator key to JAWS by preceding the accelerator key combination with Insert+F3.

8.4.3 How to Change the Look and Feel of SQL Developer

You can change the default look and feel for Oracle SQL Developer in the **Tools > Preferences > Environment** page. The look and feel determines the display colors and shapes of objects like menus and buttons.

8.4.4 How to Customize the Fonts in Code Editors

You can change the font and font size that display in editors in the **Tools > Preferences > Code Editor > Fonts** page.

8.4.5 How to Customize Syntax Highlighting

You can change the font style, as well as the foreground and background colors used in syntax highlighting within the source editor in the **Tools > Preferences > Code Editor > PL/SQL Syntax Colors** page.

8.4.6 How to Display Line Numbers in Code Editors

You can display or hide line numbers in the source editor in the **Tools > Preferences > Code Editor > Line Gutter** page.

8.4.7 How to Change the Timing for Completion Insight

You can specify the number of seconds that Completion Insight is delayed, or disable Completion Insight, in the **Tools > Preferences > Code Editor > Completion Insight** page.

8.4.8 How to Specify the Columns in the Debugger

You can choose the columns and types of information that display in the Debugger in the **Tools > Preferences > Debugger** pages.

8.5 Highly Visual Features of Oracle SQL Developer

Oracle SQL Developer includes features that are highly visual, and these features have equivalent functionality that is available to people who are blind or visually impaired:

- The UI and visual editors. The source editor provides equivalent functionality, as pages and UI elements can be completely designed and coded in the source editor.
- The Component Palette. The source editor provides equivalent functionality, as elements and tags that can be selected from the Component Palette can also be entered in the source editor.

You can add a component from the Component Palette to the UI or visual editor using keystrokes.

In Data Modeler and in the Data Miner Workflow Editor, it is possible to create, edit, and move elements on a diagram using only keystrokes.

Symbols

.trc files, [1-40](#)

A

abstract syntax tree (AST), [1-111](#)

accelerator (shortcut) keys
for menus, [1-4](#)

Active Session History (ASH) reports, [1-85](#)

advanced formatting options, [1-97](#)

Advanced Security for JDBC connection, [1-31](#)

aging policy
defined, [7-10](#)
LRU, [7-10](#)
time-based, [7-11](#)

analyzing tables, [1-23](#)

application
deploying (Application Express), [6-79](#)
importing (Application Express), [6-79](#)

Application Express
applications, [1-14](#)
deploying an application, [6-79](#)
importing an application, [6-79](#)
Publish to Apex option, [1-34](#)
refactor in bulk command, [1-14](#)
starting and stopping EPG (embedded PL/SQL gateway), [1-14](#)

applications
Application Express, [1-14](#)

Apply Filter
to display of objects, [1-3](#), [6-91](#)

ASH (Active Session History) reports, [1-85](#)

associations
file types, [6-5](#)

AST (abstract syntax tree), [1-111](#)

authentication
Kerberos
Database Advanced Parameters, [1-99](#)
OS (operating system), [6-12](#)

auto-detect formatter settings, [1-97](#)

auto-enabling RESTful services, [6-126](#)

Auto-Generated connections folder, [1-29](#)

autocommit
preference for database connections, [1-99](#)

autocomplete preferences, [1-96](#)

autogenerated connections, [1-29](#)

Automated Workload Repository (AWR) reports, [1-85](#)

autotrace
Autotrace pane, [1-46](#)
preferences for, [1-99](#)

AWR (Automated Workload Repository) reports, [1-85](#)

B

bequeath connection type, [6-12](#)

Berkeley DB
data storage for Subversion repository, [6-142](#)

bind variables
for reports, [1-84](#)
saving to disk on exit, [1-108](#)

blogs
SQL Developer, [1-129](#)

border color associated with connection, [6-11](#)

breakpoints
creating and editing, [6-133](#)

C

cache agent
start and stop dialog, [7-6](#)

cache group
create dialog, [7-6](#)
user managed, [7-7](#)

cache groups
asynchronous writethrough, [7-7](#)
dynamic, [7-6](#)
explicitly loaded, [7-6](#)
global, [7-7](#)
local, [7-7](#)
read-only, [7-7](#)
root table, [7-8](#)
selecting the columns to cache, [7-8](#)
synchronous writethrough, [7-7](#)

cache table
aging, [7-10](#)

- cart
 - command-line interface, [1-123](#)
- Cart (SQL Developer), [1-68](#)
- case insensitivity for queries
 - setting, [2-24](#)
- CDB
 - definition, [1-19](#)
- chains (Scheduler)
 - creating and editing, [6-10](#)
- Change case as you type
 - completion insight option, [1-94](#)
- charts
 - user-defined report example, [1-90](#)
- check constraints, [6-46](#)
- Check for Updates feature, [6-3](#)
- child table
 - defining foreign keys, [7-12](#)
 - selecting columns to cache, [7-12](#)
 - specifying the table attributes, [7-12](#)
- child tables
 - defining, [7-11](#)
- CLOBtoBLOB_sqldeveloper stored procedure, [1-111](#)
- close all SQL Worksheets automatically on disconnect
 - preference for database connections, [1-99](#)
- clustered primary keys (Sybase)
 - generating as Oracle index-organized tables, [1-111](#)
- clustered unique indexes (Sybase)
 - generating as Oracle index-organized tables, [1-111](#)
- coalescing an index, [1-17](#)
- code coverage
 - gathering and viewing statistics, [4-6](#)
- code fragments, [1-50](#)
- code insight (completion insight), [1-96](#)
- column sequences, [6-48](#)
- columns
 - decrypting, [1-23](#)
 - encrypting, [1-23](#)
 - normalizing column data, [1-23](#)
- command line interface to SQL Developer, [1-123](#)
- command-line interface, [1-123](#)
- Compact option for shrinking a table, [1-23](#)
- compiling
 - function
 - for debug, [1-16](#), [1-20](#)
 - view, [1-26](#)
- completion
 - SQL Developer preferences, [1-96](#)
- completion insight, [1-96](#)
- compound triggers, [1-115](#)
- configuring
 - third-party databases, [2-14](#)
- configuring file type associations, [6-5](#)
- connection
 - new, [7-2](#)
- connection color, [6-11](#)

- connections
 - creating, editing, or selecting, [6-11](#)
 - explanation, [1-28](#)
 - Hive, [6-15](#)
 - IBM DB2, [6-15](#)
 - JDBC, [6-16](#)
 - Microsoft SQL Server, [6-17](#)
 - MySQL, [6-16](#)
 - Oracle Database, [6-12](#)
 - Oracle TimesTen, [6-14](#)
 - PostgreSQL, [6-16](#)
 - separate unshared connection preference, [1-108](#)
 - separate unshared worksheet for, [1-30](#), [1-41](#)
 - Sybase Adaptive Server, [6-17](#)
 - Teradata, [6-17](#)
 - using folders to group, [1-31](#)
- constraint
 - dropping, [7-20](#)
- constraints
 - check, [6-46](#)
 - disabled, [1-87](#)
 - unique, [6-45](#)
- converted model
 - correcting errors in, [2-19](#)
 - creating and customizing, [2-18](#)
- Create Local Connections command, [1-29](#)
- Create Unit Test wizard, [6-152](#)
- create users option (migration generation options), [1-111](#)
- credentials (Scheduler)
 - creating and editing, [6-10](#)
- crossedition triggers, [1-25](#)
- customizing SQL Developer
 - setting preferences, [1-92](#)
- CVS (Concurrent Versions System)
 - preferences for, [1-118](#)
 - SQL Developer support for, [1-53](#)

D

- data
 - entering and modifying, [1-34](#)
- Data Modeler
 - included in SQL Developer, [1-124](#)
- data move
 - number of parallel streams, [1-112](#)
 - options, [1-111](#)
- Data Pump Export wizard, [6-70](#)
- Data Pump Import wizard, [6-72](#)
- data redaction, [1-23](#)
- data types
 - creating, [6-59](#)
- database administrator (DBA) features in SQL Developer, [1-54](#)
- database connections
 - creating, editing, or selecting, [6-11](#)
 - explanation, [1-28](#)

- database connections (*continued*)
 - Hive, [6-15](#)
 - IBM DB2, [6-15](#)
 - JDBC, [6-16](#)
 - Microsoft SQL Server, [6-17](#)
 - MySQL, [6-16](#)
 - Oracle Database, [6-12](#)
 - Oracle TimesTen, [6-14](#)
 - PostgreSQL, [6-16](#)
 - Sybase Adaptive Server, [6-17](#)
 - Teradata, [6-17](#)
- database destination groups (Scheduler)
 - creating and editing, [6-27](#)
- database destinations (Scheduler)
 - creating and editing, [6-27](#)
- database instance viewer
 - preferences, [1-100](#)
- database links
 - creating or editing, [6-28](#)
- database objects
 - exporting, [6-77](#)
 - finding (searching for), [1-52](#)
 - recent, [1-53](#)
 - unloading, [6-77](#)
- database schema differences
 - differences between two database schemas, [6-75](#), [6-82](#)
- DB2
 - before migrating, [2-14](#)
- DB2 (IBM) connections, [6-15](#)
- DBA features in SQL Developer
 - using, [1-54](#)
- DBMS_OUTPUT
 - user-defined report example, [1-91](#)
- DBMS_OUTPUT pane, [1-46](#)
- debugging PL/SQL function or procedure
 - dialog box, [6-132](#)
 - remote debugging, [1-39](#)
- decryption
 - data in a table column, [1-23](#)
- default password
 - Use Default Password connection option, [6-17](#)
- deploying objects using the Cart, [1-68](#)
- deployment
 - after database migration, [2-29](#)
- destination groups (Scheduler)
 - creating and editing, [6-27](#)
- destination schema
 - database schema differences, [6-75](#), [6-82](#)
- destinations
 - database (Scheduler), [1-15](#)
- dialog boxes and tabs, [6-1](#)
- diff
 - command-line interface, [1-123](#)
- directories (directory objects), [1-16](#)
- disabled constraints, [1-87](#)
- discussion forum

- discussion forum (*continued*)
 - SQL Developer, [1-129](#)
- documentation
 - generating for schema, [1-30](#)
- downloads
 - SQL Developer, [1-129](#)
- drag and drop effects, [1-99](#)
- drivers
 - third-party JDBC, [1-99](#)
- dynamic HTML
 - user-defined report example, [1-91](#)

E

- editions
 - crossedition triggers, [1-25](#)
- embedded PL/SQL gateway (EPG) for Application Express
 - starting and stopping, [1-14](#)
- emulation user
 - generating separate, [1-111](#)
- Enable Semantic Analysis Info Tip, [1-94](#)
- encryption
 - data in a table column, [1-23](#)
- entering data in tables, [1-34](#)
- EPG (embedded PL/SQL gateway) for Application Express
 - starting and stopping, [1-14](#)
- errors
 - parse exceptions, [2-19](#)
- Excel
 - importing data, [6-24](#)
- execution plan, [1-46](#)
- expired password
 - resetting, [1-30](#)
- EXPLAIN PLAN
 - execution plan, [1-46](#)
- export
 - command-line interface, [1-123](#)
- exporting
 - database objects, [6-77](#)
 - reports, [1-83](#)
 - table data, [1-23](#), [6-85](#)
 - unit test objects, [4-12](#)
- expression watches, [1-40](#)
- extensions
 - loading, [6-110](#)
 - SQL Developer, [6-112](#)
 - user-defined, [1-99](#)
- external destination groups (Scheduler), [1-15](#)
- external editor
 - SQL Developer preferences, [1-110](#)
- external tables
 - properties, [6-52](#)
- external tools, [1-11](#), [6-87](#), [6-88](#)

F

- F10 key
 - for File menu, [1-4](#)
- F4 key
 - Alt+F4 to close window with focus, [1-4](#)
- failed objects
 - generating, [1-111](#)
- features
 - managing, [6-112](#)
- file types
 - associating with SQL Developer, [6-5](#)
- file watchers (Scheduler)
 - creating and editing, [6-30](#)
- file-oriented development
 - SQL Worksheet right-click operations, [1-42](#)
- Files navigator, [1-5](#)
- filter
 - applying, [1-3](#), [6-91](#)
 - clearing, [1-3](#), [6-91](#)
- finding database objects, [1-52](#)
- Flashback Table support, [1-25](#)
- folders
 - for user-defined reports, [1-90](#), [6-62](#)
 - in Connections navigator, [1-31](#)
- font
 - increasing or decreasing size for help text, [1-125](#)
- foreign key constraint
 - creating, [7-19](#)
- foreign keys, [6-46](#)
- formatter
 - command-line interface, [1-48](#), [1-123](#)
- Freeze View
 - pinning an object's display, [1-6](#), [1-34](#), [1-83](#)
- FROM clause, [6-66](#)
- full pathname for java.exe, [1-3](#)
- functions
 - compiling, [1-16](#), [1-20](#)
 - compiling for debug, [1-16](#), [1-20](#)
 - creating, [6-31](#)
 - debugging, [6-132](#)
 - running, [6-132](#)

G

- Gather Code Coverage Statistics option, [4-6](#)
- Gather Schema Statistics menu command, [1-30](#)
- gauges
 - in SQL Worksheet and user-defined reports, [1-49](#)
- Generate DB Doc menu command, [1-30](#)
- generate failed objects option, [1-111](#)
- generate separate emulation user option, [1-111](#)
- Git
 - preferences for, [1-117](#)
- grant information
 - not migrated to Oracle, [2-12](#)
- graphical user interface (GUI)

- graphical user interface (GUI) (*continued*)
 - unit testing, [4-3](#)
- GROUP BY clause, [6-67](#)
- groups
 - connections in folders, [1-31](#)

H

- HAVING clause, [6-67](#)
- help
 - increasing or decreasing help text font size, [1-125](#)
 - using the online help, [1-125](#)
- hierarchical profiler (PL/SQL), [1-40](#)
- Hive connections, [6-15](#)
- HTML
 - dynamic (user-defined report), [1-91](#)

I

- IBM DB2
 - before migrating, [2-14](#)
- IBM DB2 connections, [6-15](#)
- IDE_USER_DIR location, [1-122](#)
- IM column store, [6-47](#)
- implementation
 - adding for a unit test, [6-152](#)
 - renaming for a unit test, [6-155](#)
- import
 - command-line interface, [1-123](#)
- importing
 - Microsoft Excel data, [6-24](#)
 - reports, [1-83](#)
 - unit test objects, [4-12](#)
- In-Memory Column Store, [6-47](#)
- index-organized tables
 - properties, [6-52](#)
- indexes
 - coalescing, [1-17](#)
 - creating or editing, [6-28](#)
 - explanation, [1-17](#)
 - making unusable, [1-17](#)
 - rebuilding, [1-17](#)
- instance viewer
 - preferences, [1-100](#)
- introduction
 - SQL Developer, [1-1](#)

J

- Java sources, [1-18](#)
- java.exe
 - pathname for, [1-3](#)
- JDBC connections, [6-16](#)
- JDBC drivers, [1-99](#)
- JDK location
 - specified in product.conf, [1-123](#)
- job classes (Scheduler)

- job classes (Scheduler) (*continued*)
 - creating and editing, [6-34](#)
- job scheduler
 - using in SQL Developer, [1-67](#)
- jobs
 - scheduling using SQL Developer, [1-67](#)
- jobs (Scheduler)
 - creating and editing, [6-32](#)
- jTDS
 - JDBC driver for Microsoft SQL Server and Sybase Adaptive Server, [1-103](#)

K

- Kerberos authentication
 - Database Advanced Parameters, [1-99](#)
- keyboard shortcuts, [1-115](#), [1-116](#)

L

- least privilege schema migration, [1-111](#)
- library
 - adding an item for unit testing, [6-151](#)
 - unit testing, [4-11](#)
- link
 - database, [6-28](#)
- Load Extension, [6-110](#)
- LOB parameters, [6-49](#)
- local connections
 - creating, [1-29](#)
- locking a table, [1-23](#)
- lookups
 - adding for unit testing, [6-151](#)
- lowercase characters in object names
 - quotation marks required, [6-1](#)

M

- materialized view logs
 - creating and editing, [6-35](#)
- materialized views
 - options when creating or editing, [6-63](#)
- Microsoft Excel
 - importing data, [6-24](#)
- Microsoft SQL Server
 - connections, [6-17](#)
 - creating data files, [2-20](#)
 - third-party JDBC drivers, [1-99](#)
- migration
 - before you start (general), [2-12](#)
 - before you start (source-specific), [2-14](#)
 - command-line interface, [1-123](#)
 - deployment after migrating, [2-29](#)
 - introduction and main topics, [2-1](#)
 - overview, [2-6](#)
 - planning for, [2-7](#)
 - quick start, [2-1](#)

- migration (*continued*)
 - required roles and privileges, [2-13](#)
- migration plan
 - preparing, [2-7](#)
- migration repository
 - creating database user for, [2-12](#)
- MIME type
 - specifying external editor for BLOB data, [1-110](#)
- MOD_PLSQL
 - OWA output, [1-47](#)
- modifying data in tables, [1-34](#)
- Monitor Sessions command, [1-7](#)
- Monitor SQL command, [1-7](#)
- moving a table to another tablespace, [1-23](#)
- multitenant container database (CDB)
 - definition, [1-19](#)
- MySQL
 - configuring before migrating, [2-16](#)
 - creating data files, [2-20](#)
 - third-party JDBC drivers, [1-99](#)
- MySQL connections
 - zero date handling, [6-16](#)

N

- navigation tree filtering option, [1-101](#)
- new connection dialog, [7-2](#)
- New Worksheet to use unshared connection
 - preference for database SQL Worksheet using unshared connection automatically, [1-108](#)
- normalizing
 - column data, [1-23](#)

O

- OBEs (Oracle By Example lessons), [1-129](#)
- objects
 - database, [1-13](#)
 - exporting, [6-77](#)
 - unloading, [6-77](#)
- OCI driver
 - database advanced parameters preference, [1-100](#)
- offline data transfer, [2-20](#)
- OLAP DML statements
 - entering in SQL Worksheet, [1-50](#)
- online help
 - using, [1-125](#)
- Open command on Connections navigator context menu, [1-13](#)
- open SQL Worksheet automatically
 - preference for database connections, [1-108](#)
- operating system authentication, [6-12](#)
- ORA_ROWSCN pseudocolumn
 - "Use ORA_ROWSCN for DataEditor insert and update statements" preference, [1-99](#)
- Oracle Application Express

- Oracle Application Express (*continued*)
 - applications, [1-14](#)
- Oracle By Example (OBE) lessons, [1-129](#)
- Oracle Database 12c features
 - generating in migration, [1-111](#)
- Oracle Database connections, [6-12](#)
- Oracle Web Agent (OWA), [1-47](#)
- ORDER BY clause, [6-67](#)
- OS authentication, [6-12](#)
- OTN page
 - SQL Developer, [1-129](#)
- overview
 - SQL Developer, [1-1](#)
- overview display of database objects
 - Open command on context menu, [1-13](#)
- OWA_OUTPUT pane, [1-47](#)

P

- packages
 - creating, [6-36](#)
 - debugging, [6-132](#)
 - running, [6-132](#)
- parallel data move streams, [1-112](#)
- parameter insight, [1-96](#)
- parse exception errors, [2-19](#)
- partitioned tables
 - partition definitions, [6-51](#)
 - partitioning options, [6-50](#)
 - subpartitioning options (templates), [6-51](#)
- password
 - resetting expired, [1-30](#)
- pathname
 - for java.exe, [1-3](#)
- PDB
 - definition, [1-19](#)
- Perforce
 - preferences for, [1-120](#)
- pinning an object's display, [1-6](#), [1-34](#), [1-83](#)
- PL/SQL hierarchical profiler, [1-40](#)
- PL/SQL packages
 - creating, [6-36](#)
- PL/SQL subprograms
 - debugging
 - breakpoints, [6-133](#)
 - running, [6-132](#)
- plsql-dbms_output
 - user-defined report example, [1-91](#)
- pluggable database (PDB)
 - definition, [1-19](#)
- pluggable databases
 - cloning, [6-5](#)
 - cloning to Oracle Cloud, [6-4](#)
 - creating, [6-36](#), [6-114](#)
 - dropping, [6-84](#)
 - plugging in, [6-120](#)
 - unplugging, [6-156](#)

- PostgreSQL
 - connections, [6-16](#)
 - third-party JDBC drivers, [1-103](#)
- preferences
 - customizing SQL Developer, [1-92](#)
- preferences (options), [1-116](#)
- primary key, [6-44](#)
- primary key attributes
 - changing, [7-18](#)
- primary key constraint
 - creating, [7-18](#)
- private database links, [1-15](#)
- private synonyms, [1-23](#), [6-41](#)
- privileges
 - required for migration, [2-13](#)
- procedures
 - creating, [6-31](#)
- product.conf file, [1-123](#)
- profiler (PL/SQL hierarchical), [1-40](#)
- profiling a PL/SQL function or procedure, [6-132](#)
- programs (Scheduler)
 - creating and editing, [6-37](#)
- projects
 - creating plan, [2-10](#)
- proxy authentication, [1-32](#)
- proxy connections, [1-32](#), [6-12](#)
- public database links, [1-15](#)
- public synonyms, [1-23](#), [6-41](#)
- Publish to Apex (Application Express) option, [1-34](#)

Q

- query builder, [1-47](#)
- query optimizer
 - statistics detail, [7-24](#)
- quotation marks
 - for name with lowercase characters, special characters, or spaces, [6-1](#)

R

- Raptor
 - former code name for SQL Developer, [1-1](#)
- rebuilding an index, [1-17](#)
- recent objects, [1-53](#)
- recompiling
 - view, [1-26](#)
- Recycle bin, [1-22](#)
- redaction, [1-23](#)
- refactor in bulk, [1-14](#)
- remote debugging, [1-39](#)
- replication schemes, [1-22](#)
- report navigator, [1-82](#)
- reports
 - bind variables for, [1-84](#)
 - exporting, [1-83](#)
 - importing, [1-83](#)

- reports (*continued*)
 - shared, [1-83](#)
 - unit testing, [4-12](#)
 - user-defined
 - chart example, [1-90](#)
 - creating and editing, [6-60](#)
 - dynamic HTML example, [1-91](#)
 - explanation, [1-89](#)
 - folders for, [1-90](#), [6-62](#)
 - gauge example, [1-49](#)
 - UserReports.xml, [1-89](#)
- repository
 - unit test, [4-4](#)
- Reset Password menu command, [1-30](#)
- resource handler
 - creating, [6-129](#)
- resource module, [3-12](#)
- resource template
 - creating, [6-128](#)
- REST Data Services administration, [3-1](#)
- RESTful service module
 - creating, [6-128](#)
- RESTful services
 - auto-enabling for schema, table, or view, [6-126](#)
- roles
 - creating and editing, [6-38](#)
 - required for migration, [2-13](#)
- root table
 - primary key, [7-9](#)
 - where clause, [7-10](#)
- root table definition
 - verifying, [7-11](#)
- root table dialog, [7-8](#)
- route pattern, [3-12](#)
- running PL/SQL function or procedure
 - dialog box, [6-132](#)

S

- scheduler
 - using in SQL Developer, [1-67](#)
- schedules (Scheduler)
 - creating and editing, [6-39](#)
- scheduling jobs in SQL Developer, [1-67](#)
- schema
 - XML, [1-28](#), [6-69](#), [6-70](#)
- Schema Browser, [1-29](#)
- schema differences
 - differences between two database schemas, [6-75](#), [6-82](#)
- schema documentation
 - generating, [1-30](#)
- schema objects, [1-13](#)
- schema statistics
 - gathering, [1-30](#)
- scratch editor, [2-34](#)
- script runner, [1-45](#)

- scripts
 - running, [1-45](#)
- search
 - for database objects, [1-52](#)
- Secure Shell (SSH)
 - authentication, [1-33](#)
 - connections, [1-33](#)
- security
 - Advanced Security for JDBC connection, [1-31](#)
- SELECT clause, [6-66](#)
- Semantic Analysis Info Tip
 - enabling, [1-94](#)
- sequences
 - creating and editing, [6-40](#)
 - for populating table columns, [6-48](#)
- SerDe class (Hive tables), [6-54](#)
- sessions
 - monitoring, [1-7](#)
- SetJavaHome line in product.conf file, [1-123](#)
- shared reports, [1-83](#)
- shared repository, [4-4](#)
- shortcut keys
 - for menus, [1-4](#)
 - restoring to default scheme, [1-115](#)
- shrinking a table, [1-23](#)
- single record view, [6-138](#)
- size of help text
 - increasing or decreasing, [1-125](#)
- snippets
 - user-defined, [1-51](#)
- source database
 - capturing, [2-17](#)
- source schema
 - database schema differences, [6-75](#), [6-82](#)
- spaces in object names
 - quotation marks required, [6-1](#)
- Spatial support
 - Maintain Spatial Metadata dialog box, [6-111](#)
 - Validate Geometry dialog box, [6-156](#)
- special characters in object names
 - quotation marks required, [6-1](#)
- split
 - data pane for a table or view, [1-36](#)
- SQL
 - monitoring execution, [1-7](#)
- SQL Developer preferences, [1-92](#)
- SQL file
 - creating, [6-40](#)
- SQL formatting
 - command-line interface, [1-48](#), [1-123](#)
- SQL scripts
 - running, [1-45](#)
- SQL Server
 - configuring, [2-15](#)
 - creating data files, [2-20](#)
 - third-party JDBC drivers, [1-99](#)
- SQL Server connections, [6-17](#)

- SQL Trace files (.trc), [1-40](#)
- SQL Worksheet
 - closing automatically on disconnect, [1-99](#)
 - entering OLAP DML statements, [1-50](#)
 - opening automatically on database connection, [1-108](#)
 - unshared for a connection, [1-30](#), [1-41](#)
 - using, [1-41](#)
- SSH
 - authentication, [1-33](#)
 - connections, [1-33](#)
- Start Page for SQL Developer, [1-129](#)
- statistics
 - computing table and column, [1-23](#)
 - estimating table and column, [1-23](#)
 - gathering schema, [1-30](#)
- storage options, [6-54](#)
- stored procedures
 - generating for Migrate Blobs Offline, [1-111](#)
- subpartitions
 - templates, [6-51](#)
- subprograms
 - creating, [6-31](#)
 - debugging, [6-132](#)
 - running, [6-132](#)
- substitution variables, [1-45](#)
- Subversion
 - preferences for, [1-117](#)
 - SQL Developer support for, [1-53](#)
- SVN
 - preferences for, [1-117](#)
- Sybase Adaptive Server
 - connections, [6-17](#)
 - creating data files, [2-20](#)
 - third-party JDBC drivers, [1-99](#)
- Sybase clustered primary keys
 - generating as Oracle index-organized tables, [1-111](#)
- Sybase clustered unique indexes
 - generating as Oracle index-organized tables, [1-111](#)
- Synchronize Unit Test wizard, [6-155](#)
- synonyms
 - creating and editing, [6-41](#)

T

- T-SQL
 - not supported in worksheet for SQL Server or Sybase connection, [2-34](#)
 - translation to PL/SQL, [2-34](#)
- tables
 - analyzing, [1-23](#)
 - compacting, [1-23](#)
 - computing statistics, [1-23](#)
 - creating and editing, [6-42](#)
 - creating quickly, [6-41](#)

- tables (*continued*)
 - decrypting column data, [1-23](#)
 - encrypting column data, [1-23](#)
 - entering and modifying data, [1-34](#)
 - estimating statistics, [1-23](#)
 - exporting data, [1-23](#), [6-85](#)
 - external
 - properties, [6-52](#)
 - in-memory, [6-47](#)
 - index-organized
 - properties, [6-52](#)
 - locking, [1-23](#)
 - moving to another tablespace, [1-23](#)
 - normalizing column data, [1-23](#)
 - partitioned
 - partition definitions, [6-51](#)
 - partitioning options, [6-50](#)
 - subpartitioning options (templates), [6-51](#)
 - shrinking, [1-23](#)
 - splitting the data pane, [1-36](#)
 - temporary, [6-42](#)
 - truncating, [1-23](#)
 - validating structure, [1-23](#)
- tablespaces
 - creating and editing, [6-55](#)
- Target Oracle Database 12c features in migration
 - option, [1-111](#)
- temporary tables, [6-42](#)
- Teradata
 - before migrating, [2-16](#)
 - connections, [6-17](#)
- test case
 - definition of, [4-2](#)
- test suite
 - definition of, [4-2](#)
- thick (OCI) driver
 - database advanced parameters preference, [1-100](#)
- third-party databases
 - configuring before migrating, [2-14](#)
- third-party JDBC drivers, [1-99](#)
- time-based aging, [7-11](#)
- TimesTen
 - cache groups, [1-14](#)
 - database connections, [6-14](#)
 - replication schemes, [1-22](#)
 - support for in SQL Developer, [1-125](#)
 - tasks (how to perform in SQL Developer), [7-1](#)
- Tip of the Day, [1-126](#)
- TKPROF
 - opening in SQL Developer as an alternative to using, [1-40](#)
- tnsnames.ora file, [1-29](#)
- tools
 - external, [1-11](#), [6-87](#), [6-88](#)
- trace files (.trc), [1-40](#)
- transferring
 - data offline, [2-20](#)

- translation scratch editor, [2-34](#)
- triggers
 - compound, [1-115](#)
 - creating and editing, [6-57](#)
- truncating a table, [1-23](#)
- tutorial
 - creating objects for a small database, [5-1](#)
- types
 - creating, [6-59](#)

U

- unique constraints, [6-45](#)
- unit test
 - definition of (as related to test cases), [4-2](#)
- unit testing
 - adding a category for lookups, [6-151](#)
 - adding an implementation, [6-152](#)
 - adding an item to a library, [6-151](#)
 - best practices, [4-14](#)
 - command-line interface, [1-123](#), [4-13](#)
 - copying or renaming a unit test, [6-152](#)
 - Create Unit Test wizard, [6-152](#)
 - editing and running a unit test, [4-5](#)
 - example, [4-16](#)
 - exporting objects, [4-12](#)
 - graphical user interface (GUI), [4-3](#)
 - importing objects, [4-12](#)
 - introduction and main topics, [4-1](#)
 - library, [4-11](#)
 - managing users, [6-155](#)
 - overview, [4-1](#)
 - renaming an implementation, [6-155](#)
 - reports, [4-12](#)
 - repository, [4-4](#)
 - Synchronize Unit Test wizard, [6-155](#)
 - tutorial, [4-16](#)
 - variable substitution, [4-10](#)
- unloading
 - database objects, [6-77](#)
- unshared connection
 - preference for using on new worksheet, [1-108](#)
- unshared worksheet for a connection, [1-30](#), [1-41](#)
- unusable indexes, [1-17](#)
- updates
 - checking for SQL Developer updates, [6-3](#)
 - managing, [6-112](#)
- URI pattern, [3-12](#)
- URI template, [3-12](#)
- usage suggestions, [1-126](#)
- Use as Template command on context menu, [1-13](#)
- Use Default Password connection option, [6-17](#)
- Use Windows Authentication connection option, [6-17](#)
- user information directory (IDE_USER_DIR), [1-122](#)
- user interface
 - restoring to original settings, [1-12](#)

- user interface (UI)
 - unit testing, [4-3](#)
- user-defined extensions, [1-99](#)
- user-defined reports
 - chart example, [1-90](#)
 - creating and editing, [6-60](#)
 - dynamic HTML example, [1-91](#)
 - explanation, [1-89](#)
 - folders for, [1-90](#), [6-62](#)
 - gauge example, [1-49](#)
 - UserReports.xml, [1-89](#)
- user-defined snippets, [1-51](#)
- user-defined types
 - creating, [6-59](#)
- UserReports.xml, [1-89](#)
- users
 - creating and editing, [6-59](#)
- using, [1-54](#)
- UT_REPO_ADMINISTRATOR role, [4-5](#)
- UT_REPO_USER role, [4-5](#)

V

- variable substitution
 - unit testing, [4-10](#)
- versioning
 - preferences for, [1-116](#)
 - SQL Developer support for, [1-53](#)
- viewlets
 - SQL Developer, [1-129](#)
- views
 - compiling, [1-26](#)
 - creating and editing, [6-62](#)
 - FROM clause, [6-66](#)
 - GROUP BY clause, [6-67](#)
 - HAVING clause, [6-67](#)
 - options when creating or editing, [6-63](#)
 - ORDER BY clause, [6-67](#)
 - recompiling, [1-26](#)
 - SELECT clause, [6-66](#)
 - splitting the data pane, [1-36](#)
 - WHERE clause, [6-67](#)

W

- watches, [1-40](#)
- WHERE clause, [6-67](#)
- window groups (Scheduler)
 - creating and editing, [6-69](#)
- windows (Scheduler)
 - creating and editing, [6-68](#)
- Windows authentication
 - Use Windows Authentication connection option, [6-17](#)

X

XML DB Repository, [1-27](#)

XML schema

 creating, [6-69](#), [6-70](#)

xUnit frameworks, [4-1](#)

Z

zero date handling

 MySQL connections, [6-16](#)