

Oracle® Fusion Middleware

Metadata Repository Builder's Guide for Oracle Business
Intelligence Enterprise Edition

11g Release 1 (11.1.1)

E10540-02

April 2011

Oracle Fusion Middleware Metadata Repository Builder's Guide for Oracle Business Intelligence Enterprise Edition, 11g Release 1 (11.1.1)

E10540-02

Copyright © 2010, 2011, Oracle and/or its affiliates. All rights reserved.

Primary Author: Marla Azriel

Contributing Author: Bob Ertl

Contributors: Oracle Business Intelligence development, product management, and quality assurance teams

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, the following notice is applicable:

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

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 licensed through X/Open Company, Ltd.

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

Contents

Preface	xxi
Audience	xxi
Documentation Accessibility	xxi
Related Documents	xxii
Conventions	xxii
New Features for Oracle BI Metadata Repository Builders	xxiii
New Features for Oracle BI EE 11g Release 1 (11.1.1.5)	xxiii
New Features for Oracle BI EE 11g Release 1 (11.1.1.3)	xxv
1 Introduction to Building Your Metadata Repository	
About Oracle BI Server and Oracle BI Repository Architecture	1-1
About Oracle BI Server Architecture	1-1
About Layers in the Oracle BI Repository	1-3
Planning Your Business Model	1-4
Analyzing Your Business Model Requirements	1-4
Identifying the Content of the Business Model	1-5
Identifying Logical Fact Tables	1-5
Identifying Logical Dimension Tables	1-6
Identifying Dimensions	1-7
About Dimensions with Multiple Hierarchies	1-8
Identifying Lookup Tables	1-8
Identifying the Data Source Content for the Physical Layer	1-8
About Types of Physical Schemas in Relational Data Sources	1-9
About Cubes in Multidimensional Data Sources	1-9
Identifying the Data Source Table Structure	1-10
Guidelines for Designing a Repository	1-10
General Tips for Working on the Repository	1-11
Design Tips for the Physical Layer	1-11
Design Tips for the Business Model and Mapping Layer	1-12
Modeling Outer Joins	1-14
Design Tips for the Presentation Layer	1-14
Topics of Interest in Other Guides	1-15
System Requirements and Certification	1-16

2 Before You Begin

About the Oracle BI Administration Tool	2-1
Opening the Administration Tool	2-2
About the Administration Tool Main Window	2-2
Setting Preferences	2-3
About Administration Tool Menus	2-6
File Menu.....	2-6
Edit Menu.....	2-8
View Menu.....	2-8
Manage Menu.....	2-8
Tools Menu	2-9
Actions Menu	2-10
Window Menu.....	2-10
Help Menu	2-10
Using the Physical and Business Model Diagrams	2-10
Editing, Deleting, and Reordering Objects in the Repository	2-13
About Naming Requirements for Repository Objects.....	2-13
Using the Browse Dialog to Browse for Objects.....	2-13
Changing Icons for Repository Objects.....	2-14
Sorting Objects in the Administration Tool	2-15
About Features and Options for Oracle Marketing Segmentation.....	2-15
About the Oracle BI Server Command-Line Utilities	2-15
Running bi-init to Launch a Shell Window Initialized to Your Oracle Instance	2-16
About Options in Fusion Middleware Control and NQSConfig.INI	2-17
About the SampleApp.rpd Demonstration Repository	2-17
Using Online and Offline Repository Modes	2-18
Opening a Repository in Offline Mode.....	2-18
Opening a Repository in Online Mode	2-19
Guidelines for Using Online Mode	2-19
Checking Out Objects	2-20
Checking In Changes.....	2-20
About Read-Only Mode.....	2-21
Checking the Consistency of a Repository or a Business Model	2-21
About the Consistency Check Manager.....	2-21
Checking the Consistency of Repository Objects	2-22
Using the validate.rpd Utility to Check Repository Consistency	2-23

3 Setting Up and Using the Multiuser Development Environment

About the Multiuser Development Environment	3-1
About the Multiuser Development Process	3-2
Setting Up Projects	3-3
About Projects.....	3-4
About the Project Dialog.....	3-4
Creating Projects.....	3-5
About Converting Older Projects During Repository Upgrade	3-7
Setting Up the Multiuser Development Directory	3-7
Identifying the Multiuser Development Directory	3-7

Copying the Master Repository to the Multiuser Development Directory	3-8
Setting Up a Pointer to the Multiuser Development Directory	3-8
Making Changes in a Multiuser Development Environment	3-8
Checking Out Repository Projects	3-9
About Repository Project Checkout	3-9
Checking Out Projects	3-10
Using the extractprojects Utility to Extract Projects	3-11
About Changing and Testing Metadata	3-12
About Multiuser Development Menu Options	3-12
About Closing a Repository Before Publishing It to the Network	3-13
Checking In Multiuser Development Repository Projects	3-13
About the Multiuser Development Merge Process	3-14
How are Multiuser Merges Different from Standard Repository Merges?	3-15
Checking In Projects	3-15
Tracking Changes to the Master Repository	3-17
Branching in Multiuser Development	3-17
About Branching	3-17
Using the Multi-Team, Multi-Release Model in Oracle Business Intelligence	3-19
Synchronizing RPD Branches	3-19
Viewing and Deleting History for Multiuser Development	3-19
Viewing Multiuser Development History	3-20
Deleting Multiuser Development History	3-21
Setting Multiuser Development Options	3-21

4 Importing Metadata and Working with Data Sources

Creating New Oracle BI Repository Files	4-1
Performing Data Source Preconfiguration Tasks	4-2
Setting Up ODBC Data Source Names (DSNs)	4-3
Setting Up Oracle Database Data Sources	4-4
Setting Up Oracle OLAP Data Sources	4-4
Setting Up Oracle TimesTen In-Memory Database Data Sources	4-4
Avoiding Query Failures Due to Limited Connections with TimesTen Data Sources	4-5
Setting Up Essbase Data Sources	4-6
Updating Essbase Information in opmn.xml	4-6
Adding Essbase Information to bi-init.cmd	4-7
Setting Up Hyperion Financial Management Data Sources	4-7
Setting Up SAP/BW Data Sources	4-10
Setting Up Oracle RPAS Data Sources	4-10
Importing Metadata from Relational Data Sources	4-10
About the Map to Logical Model and Publish to Warehouse Screens	4-13
Importing Metadata from Multidimensional Data Sources	4-14
About Importing Metadata from Oracle RPAS Data Sources	4-18
Importing Metadata from XML Data Sources	4-19
About Using XML as a Data Source	4-19
Importing Metadata from XML Data Sources Using the XML Gateway	4-20
Examples of XML Documents Generated by the Oracle BI Server XML Gateway	4-23
Accessing HTML Tables	4-28

Importing Metadata from XML Data Sources Using XML ODBC	4-29
Example of an XML ODBC Data Source	4-30
Examples of XML Documents	4-31
Using a Standby Database with Oracle Business Intelligence	4-34
About Using a Standby Database with Oracle Business Intelligence	4-34
Creating the Database Object for the Standby Database Configuration.....	4-35
Creating Connection Pools for the Standby Database Configuration.....	4-36
Updating Write-Back Scripts in a Standby Database Configuration.....	4-38
Setting Up Usage Tracking in a Standby Database Configuration.....	4-38
Setting Up Event Polling in a Standby Database Configuration.....	4-39
Setting Up Oracle BI Scheduler in a Standby Database Configuration	4-39

5 Working with ADF Business Component Data Sources

What Are ADF Business Components?	5-1
About Operational Reporting with ADF Business Components	5-2
What Happens During Import?	5-2
About Specifying a SQL Bypass Database.....	5-3
Setting Up ADF Business Component Data Sources	5-3
Creating a WebLogic Domain	5-4
Deploying OBIEEBroker as a Shared Library in Oracle WebLogic Server.....	5-4
Deploying the Application EAR File to Oracle WebLogic Server from JDeveloper	5-5
Setting Up a JDBC Data Source in the WebLogic Server	5-8
Setting the Logging Level for the Deployed Application in Oracle WebLogic Server	5-8
Importing Metadata from ADF Business Component Data Sources.....	5-9
Enabling the Ability to Pass Custom Parameters to the ADF Application.....	5-11
Propagating Labels and Tooltips from ADF Business Component Data Sources.....	5-11
What are Labels and Tooltips?	5-12
About the Session Variable Naming Scheme for UI Hints	5-12
About Determining the Physical Column for a Presentation Column	5-13
Initializing Session Variables Automatically for Propagating UI Hints	5-14
Example of Using UI Hints From an Oracle ADF Data Source When Creating Analyses...	5-14
Using XML Code in Initialization Blocks to Query UI Hints	5-15

6 Setting Up Database Objects and Connection Pools

Setting Up Database Objects	6-1
About Database Types in the Physical Layer.....	6-1
Creating a Database Object Manually in the Physical Layer.....	6-2
When to Allow Direct Database Requests by Default.....	6-3
Specifying SQL Features Supported by a Data Source.....	6-4
About Connection Pools	6-5
About Connection Pools for Initialization Blocks	6-6
Creating or Changing Connection Pools	6-6
Setting Connection Pool Properties in the General Tab	6-7
Common Connection Pool Properties in the General Tab.....	6-7
Multidimensional Connection Pool Properties in the General Tab.....	6-12
Setting Connection Pool Properties in the Connection Scripts Tab.....	6-14
Setting Connection Pool Properties in the XML Tab	6-16

Setting Connection Pool Properties in the Write Back Tab	6-17
Setting Connection Pool Properties in the Miscellaneous Tab	6-19
Setting Up Persist Connection Pools	6-21
About Setting the Buffer Size and Transaction Boundary	6-22

7 Working with Physical Tables, Cubes, and Joins

Working with the Physical Diagram	7-2
Creating Physical Layer Folders	7-4
Creating Physical Layer Catalogs and Schemas	7-4
Creating Catalogs	7-4
Creating Schemas	7-5
Using a Variable to Specify the Name of a Catalog or Schema	7-5
Setting Up Display Folders in the Physical Layer	7-5
Working with Physical Tables	7-6
About Tables in the Physical Layer	7-6
About Physical Alias Tables	7-8
Creating and Managing Physical Tables and Physical Cube Tables	7-10
Creating or Editing Physical Tables	7-10
Creating Alias Tables	7-12
Setting Physical Table Properties for XML Data Sources	7-12
About the Calc Scripts Tab for Essbase Data Sources	7-12
Creating and Managing Columns and Keys for Relational and Cube Tables	7-12
Creating and Editing a Column in a Physical Table	7-13
Specifying a Primary Key for a Physical Table	7-14
Deleting Physical Columns for All Data Sources	7-14
Viewing Data in Physical Tables or Columns	7-14
Working with Multidimensional Sources in the Physical Layer	7-15
About Physical Cube Tables	7-15
About Measures in Multidimensional Data Sources	7-15
About Externally Aggregated Measures	7-16
Working with Physical Dimensions and Physical Hierarchies	7-16
Working with Physical Dimension Objects	7-17
Working with Physical Hierarchy Objects	7-17
Adding or Removing Cube Columns in a Hierarchy	7-18
Working with Cube Variables for SAP/BW Data Sources	7-19
Viewing Members in Physical Cube Tables	7-20
Working with Essbase Data Sources	7-21
About Using Essbase Data Sources with Oracle Business Intelligence	7-21
About Incremental Import	7-23
Working with Essbase Alias Tables	7-24
Determining the Value to Use for Display	7-24
Explicitly Defining Columns for Each Alias	7-24
Modeling User-Defined Attributes	7-25
Associating Member Attributes to Dimensions and Levels	7-25
Modeling Alternate Hierarchies	7-25
Modeling Measure Hierarchies	7-27
Improving Performance by Using Unqualified Member Names	7-27

Working with Hyperion Financial Management Data Sources	7-28
About Query Support for Hyperion Financial Management Data Sources	7-29
Working with Oracle OLAP Data Sources	7-30
About Importing Metadata from Oracle OLAP Data Sources	7-30
Working with Oracle OLAP Analytic Workspace (AW) Objects.....	7-31
Working with Oracle OLAP Dimensions, Hierarchies, and Levels	7-31
Working with Oracle OLAP Cubes and Columns	7-33
Working with Physical Foreign Keys and Joins	7-33
About Physical Joins	7-34
About Primary Key and Foreign Key Relationships	7-34
About Complex Joins	7-34
About Multi-Database Joins	7-35
About Fragmented Data	7-35
Defining Physical Joins with the Physical Diagram.....	7-36
Defining Physical Joins with the Joins Manager	7-37
Deploying Opaque Views	7-38
About Deploying Opaque Views.....	7-38
Deploying Opaque View Objects.....	7-38
Using the Create View SELECT Statement	7-39
Undeploying a Deployed View.....	7-40
When to Delete Opaque Views or Deployed Views	7-41
When to Redeploy Opaque Views.....	7-41
Using Hints	7-41
How to Use Oracle Hints	7-42
About the Index Hint	7-42
About the Leading Hint	7-42
About Performance Considerations for Hints	7-42
Creating Hints.....	7-42
Displaying and Updating Row Counts for Physical Tables and Columns	7-43

8 Working with Logical Tables, Joins, and Columns

Creating the Business Model and Mapping Layer	8-1
Creating Business Models.....	8-2
Automatically Creating Business Model Objects	8-2
Automatically Creating Business Model Objects for Multidimensional Data Sources	8-2
Duplicating a Business Model and Subject Area.....	8-3
Working with the Business Model Diagram	8-3
Creating and Managing Logical Tables	8-4
Creating Logical Tables.....	8-5
Creating and Managing Logical Table Sources	8-6
Specifying a Primary Key in a Logical Table	8-6
Reviewing Foreign Keys for a Logical Table.....	8-6
Defining Logical Joins	8-6
Defining Logical Joins with the Business Model Diagram	8-7
Defining Logical Joins with the Joins Manager	8-8
Creating Logical Joins with the Joins Manager	8-8
Creating Logical Foreign Key Joins with the Joins Manager.....	8-9

Specifying a Driving Table	8-10
Identifying Physical Tables That Map to Logical Objects	8-11
Creating and Managing Logical Columns	8-11
Creating Logical Columns	8-12
Basing the Sort for a Logical Column on a Different Column.....	8-12
Enabling Double Column Support by Assigning a Descriptor ID Column	8-13
Creating Derived Columns.....	8-13
Configuring Logical Columns for Multicurrency Support.....	8-14
Setting Default Levels of Aggregation for Measure Columns	8-15
Setting Up Dimension-Specific Aggregate Rules for Logical Columns.....	8-16
Defining Aggregation Rules for Multidimensional Data Sources	8-18
Associating an Attribute with a Logical Level in Dimension Tables	8-19
Moving or Copying Logical Columns	8-20
Enabling Write Back On Columns	8-21
Setting Up Display Folders in the Business Model and Mapping Layer	8-22
Modeling Bridge Tables	8-23
Creating Joins in the Physical Layer for Bridge and Associated Dimension Tables.....	8-24
Modeling the Associated Dimension Tables in a Single Dimension	8-24
Modeling the Associated Dimension Tables in Separate Dimensions	8-25

9 Working with Logical Dimensions

Creating and Managing Dimensions with Level-Based Hierarchies	9-2
About Level-Based Hierarchies	9-2
Using Dimension Hierarchy Levels in Level-Based Hierarchies.....	9-4
Manually Creating Dimensions, Levels, and Keys with Level-Based Hierarchies	9-4
Creating Dimensions in Level-Based Hierarchies.....	9-5
Creating Logical Levels in a Dimension.....	9-5
Associating a Logical Column and Its Table with a Dimension Level	9-6
Identifying the Primary Key for a Dimension Level	9-9
Selecting and Sorting Chronological Keys in a Time Dimension	9-10
Adding a Dimension Level to the Preferred Drill Path.....	9-10
Automatically Creating Dimensions with Level-Based Hierarchies.....	9-10
Populating Logical Level Counts Automatically	9-12
Creating and Managing Dimensions with Parent-Child Hierarchies	9-13
About Parent-Child Hierarchies	9-13
About Levels and Distances in Parent-Child Hierarchies	9-14
About Parent-Child Relationship Tables.....	9-15
About Parent-Child Hierarchies Populated with Preaggregated Data	9-16
Creating Dimensions with Parent-Child Hierarchies.....	9-17
Defining Parent-Child Relationship Tables.....	9-18
Adding the Parent-Child Relationship Table to the Model.....	9-20
Maintaining Parent-Child Hierarchies Based on Relational Tables	9-21
Modeling Time Series Data	9-21
About Time Series Functions.....	9-21
About the AGO Function.....	9-22
About the TODATE Function	9-23
About the PERIODROLLING Function	9-24

Creating Logical Time Dimensions	9-25
Selecting the Time Option in the Logical Dimension Dialog	9-26
Setting Chronological Keys for Each Level	9-26
Creating AGO, TODATE, and PERIODROLLING Measures	9-27

10 Managing Logical Table Sources (Mappings)

Creating Logical Table Sources	10-1
Setting Priority Group Numbers for Logical Table Sources	10-2
Defining Physical to Logical Table Source Mappings and Creating Calculated Items	10-4
Unmapping a Logical Column from Its Source	10-6
Defining Content of Logical Table Sources	10-6
Verifying that Joins Exist from Dimension Tables to Fact Table.....	10-7
About WHERE Clause Filters	10-10
Working with Parent-Child Settings in the Logical Table Source	10-10
Setting Up Aggregate Navigation by Creating Sources for Aggregated Fact Data	10-11
Setting Up Fragmentation Content for Aggregate Navigation	10-12
Specifying Fragmentation Content for Single Column, Value-Based Predicates.....	10-12
Specifying Fragmentation Content for Single Column, Range-Based Predicates	10-12
Specifying Multicolumn Content Descriptions	10-13
Specifying Parallel Content Descriptions	10-13
Examples of Parallel Content Descriptions.....	10-14
Specifying Unbalanced Parallel Content Descriptions.....	10-15
Specifying Fragmentation Content for Aggregate Table Fragments.....	10-15
Specifying the Aggregate Table Content.....	10-16
Defining a Physical Layer Table with a Select Statement to Complete the Domain....	10-17
Specifying the SQL Virtual Table Content	10-17
Creating Physical Joins for the Virtual Table.....	10-17

11 Creating and Maintaining the Presentation Layer

Creating and Customizing the Presentation Layer	11-1
Creating Subject Areas	11-2
Automatically Creating Subject Areas Based on Logical Stars and Snowflakes	11-2
Removing Unneeded or Unwanted Columns	11-3
Renaming Presentation Columns to User-Friendly Names.....	11-4
Exporting Logical Keys in the Subject Area	11-4
Setting an Implicit Fact Column in the Subject Area	11-4
Maintaining the Presentation Layer	11-4
Working with Subject Areas	11-5
Working with Presentation Tables and Columns	11-6
Creating and Managing Presentation Tables	11-6
Nesting Folders in Answers	11-7
Creating and Managing Presentation Columns	11-7
Working with Presentation Hierarchies and Levels	11-8
Creating and Managing Presentation Hierarchies	11-9
Modeling Dimensions with Multiple Hierarchies in the Presentation Layer	11-10
Editing Presentation Hierarchy Objects	11-11
Creating and Managing Presentation Levels	11-11

Setting Permissions for Presentation Layer Objects	11-12
Generating a Permission Report for Presentation Layer Objects.....	11-14
Sorting Columns in the Permissions Dialog	11-14
Creating Aliases (Synonyms) for Presentation Layer Objects	11-14
12 Creating and Persisting Aggregates for Oracle BI Server Queries	
About Aggregate Persistence in Oracle Business Intelligence	12-1
Identifying Query Candidates for Aggregation	12-2
Using the Aggregate Persistence Wizard to Generate the Aggregate Specification	12-3
Writing the Create Aggregates Specification Manually	12-6
What Constraints Are Imposed During the Create Process?.....	12-7
How to Write the Create Aggregates Specification	12-8
Adding Surrogate Keys to Dimension Aggregate Tables	12-8
About the Create/Prepare Aggregates Syntax.....	12-9
About Surrogate Key Output from Create/Prepare Aggregates	12-9
Running the Aggregate Specification Against the Oracle BI Server	12-10
Troubleshooting Aggregate Persistence	12-10
13 Applying Data Access Security to Repository Objects	
About Data Access Security	13-1
Where Do I Find Information About Security Tasks?	13-2
Setting Up Row-Level Security	13-3
Setting Up Row-Level Security (Data Filters) in the Repository.....	13-3
Setting Up Row-Level Security in the Database.....	13-6
Setting Up Object Permissions	13-8
About Permission Inheritance for Users and Application Roles	13-10
Setting Query Limits	13-12
Accessing the Query Limits Functionality in the Administration Tool.....	13-12
Limiting Queries By the Number of Rows Received.....	13-12
Limiting Queries By Maximum Run Time and Restricting to Particular Time Periods.....	13-13
Allowing or Disallowing Direct Database Requests.....	13-13
Allowing or Disallowing the Populate Privilege.....	13-14
About Applying Data Access Security in Offline Mode	13-14
Setting Up Placeholder Application Roles for Offline Repository Development	13-15
About the List of Users in the Administration Tool	13-16
14 Completing Oracle BI Repository Setup	
Configuring the Repository for Oracle Scorecard and Strategy Management	14-1
Saving the Repository and Checking Consistency	14-2
Testing and Refining the Repository	14-3
Making the Repository Available for Queries	14-5
Creating Data Source Connections to the Oracle BI Server for Client Applications	14-5
Publishing to the User Community	14-5

15 Setting Up Data Sources on Linux and UNIX

About Setting Up Data Sources on Linux and UNIX.....	15-1
Configuring Data Source Connections Using Native Gateways.....	15-2
Troubleshooting OCI Connections.....	15-4
About Updating Row Counts in Native Databases.....	15-5
Using DataDirect Connect ODBC Drivers on Linux and UNIX.....	15-6
Configuring the DataDirect Connect ODBC Driver for Microsoft SQL Server Database....	15-6
Configuring the DataDirect Connect ODBC Driver for Sybase ASE Database.....	15-8
Configuring the DataDirect Connect ODBC Driver for Informix Database.....	15-9
Configuring Database Connections Using Native ODBC Drivers.....	15-10
Configuring Oracle RPAS ODBC Data Sources on AIX UNIX.....	15-11
Configuring Essbase Data Sources on Linux and UNIX.....	15-12
Configuring DB2 Connect on IBM z/OS and s/390 Platforms.....	15-13

16 Managing Oracle BI Repository Files

Comparing Repositories.....	16-1
Turning Off Compare Mode.....	16-3
Equalizing Objects.....	16-3
About Equalizing Objects.....	16-3
Using the Equalize Objects Dialog.....	16-4
Using the equalizerpds Utility.....	16-5
About Values for TypeName.....	16-6
Merging Repositories.....	16-8
Performing Full Repository Merges.....	16-8
About Full Repository Merges.....	16-8
Performing Full Repository Merges With a Common Parent.....	16-10
Performing Full Repository Merges Without a Common Parent.....	16-15
Performing Patch Merges.....	16-16
About Patch Merges.....	16-16
Generating a Repository Patch.....	16-18
Applying a Repository Patch.....	16-19
Using patchrpd to Apply a Patch.....	16-20
Querying and Managing Repository Metadata.....	16-21
Querying the Repository.....	16-21
Constructing a Filter for Query Results.....	16-23
Querying Related Objects.....	16-25
Changing the Repository Password.....	16-26

17 Using Expression Builder and Other Utilities

Using Expression Builder.....	17-1
About the Expression Builder Dialogs.....	17-1
About the Expression Builder Toolbar.....	17-3
About the Categories in the Category Pane.....	17-4
Setting Up an Expression.....	17-5
Navigating Within Expression Builder.....	17-6
Building an Expression.....	17-6

About the INDEXCOL Conversion Function	17-7
Using Administration Tool Utilities	17-7
Using the Replace Column or Table Wizard.....	17-7
Using the Oracle BI Event Tables Utility	17-10
Using the Externalize Strings Utility	17-10
Using the Rename Wizard	17-11
Using the Update Physical Layer Wizard	17-12
Generating Documentation of Repository Mappings	17-13
Generating a Metadata Dictionary	17-14
Removing Unused Physical Objects.....	17-15
Persisting Aggregates	17-15
Using the Calculation Wizard	17-15

18 Using Variables in the Oracle BI Repository

About Repository Variables	18-1
About Static Repository Variables.....	18-2
About Dynamic Repository Variables	18-2
Creating Repository Variables	18-3
Using Repository Variables in Expression Builder	18-3
About Session Variables	18-4
About System Session Variables	18-4
About Nonsystem Session Variables.....	18-6
Creating Session Variables	18-6
Working with Initialization Blocks	18-7
About Using Initialization Blocks with Variables	18-7
Initializing Dynamic Repository Variables	18-8
Initializing Session Variables	18-8
About Row-Wise Initialization	18-8
Initializing a Variable with a List of Values.....	18-9
Creating Initialization Blocks	18-10
Assigning a Name and Schedule to Initialization Blocks	18-10
Selecting and Testing the Data Source and Connection Pool.....	18-11
Examples of Initialization Strings.....	18-13
Testing Initialization Blocks	18-14
Associating Variables with Initialization Blocks	18-15
Establishing Execution Precedence	18-16
When Execution of Session Variable Initialization Blocks Cannot Be Deferred.....	18-17
Enabling and Disabling Initialization Blocks.....	18-17

A Managing the Repository Lifecycle in a Multiuser Development Environment

Planning Your Multiuser Development Deployment	A-1
About Business Organization and Governance Process Best Practices	A-2
About Technical Team Roles and Responsibilities	A-2
Multiuser Development Architecture	A-4
About Multiuser Development Concepts.....	A-4
About Multiuser Development Styles.....	A-6

Multiuser Development Sandbox Architecture.....	A-9
Multiuser Development and Lifecycle Management Architecture	A-11
Understanding the Multiuser Development Environment	A-13
About Multiuser Development Environment Task Flow	A-14
About Multiuser Development Projects	A-14
How to Create Branches.....	A-16
How to Create a Main Branch.....	A-16
How to Create a Side Branch.....	A-16
How to Create a Delegated Administration Branch.....	A-18
Which Merge Utility Should I Use?.....	A-18
MUD Tips and Best Practices.....	A-19
Best Practices for Branching	A-20
Best Practices for Setting Up Projects.....	A-20
Best Practices for Three-Way Merges.....	A-21
Best Practices for MUD Merges	A-21
Best Practices for Two-Way Merges	A-22
Best Practices for Production Migration.....	A-23
Best Practices for Application Roles and Users	A-24
Troubleshooting Multiuser Development.....	A-24

B MUD Case Study: Eden Corporation

About the Eden Corporation Fictional Case Study	B-1
Phase I - Initiating Multiuser Development (MUD).....	B-3
Starting Initiative S.....	B-4
Setting Up MUD Projects	B-5
First Developer Checks Out.....	B-6
Second Developer Checks Out.....	B-8
First Developer Checks In.....	B-9
Second Developer Checks In	B-10
MUD Administrator Test Migration Activities	B-10
Phase I Testing.....	B-11
Phase I Migration to Production.....	B-11
Phase I Summary.....	B-12
Phase II - Branching, Fixing, and Patching	B-12
Setting Up the Second Branch.....	B-13
Developers Check Out Projects.....	B-13
Patch Fix for the Main Branch	B-13
Finishing and Merging Phase II Branch.....	B-16
Phase II Summary	B-16
Phase III - Independent Semantic Model Development.....	B-17
Security Considerations for Multiple Independent Semantic Models.....	B-17
Sales Semantic Model Developers Check Out	B-18
HR Semantic Model Developer Builds Content	B-18
Phase III Summary	B-19

C Logical SQL Reference

SQL Syntax and Semantics.....	C-2
--------------------------------------	------------

Syntax and Usage Notes for the SELECT Statement	C-2
Basic Syntax for the SELECT Statement	C-3
Usage Notes	C-3
Subquery Support	C-3
SELECT List Syntax	C-4
FROM Clause Syntax.....	C-4
WHERE Clause Syntax.....	C-4
GROUP BY Clause Syntax	C-5
ORDER BY Clause Syntax	C-5
Syntax and Usage Notes for SELECT_PHYSICAL	C-5
Syntax for the SELECT_PHYSICAL Statement	C-6
Aggregate Functions Not Supported in SELECT_PHYSICAL Queries.....	C-6
Queries Supported by SELECT_PHYSICAL	C-7
Using the NATURAL_JOIN Keyword	C-8
Special Usages of SELECT_PHYSICAL.....	C-9
Rules for Queries with Aggregate Functions.....	C-9
Computing Aggregates of Baseline Columns.....	C-9
Computing Aggregates of Measure Columns	C-11
Display Function Reset Behavior.....	C-12
Alternative Syntax	C-13
Using FILTER to Compute a Conditional Aggregate.....	C-13
Operators.....	C-14
SQL Logical Operators	C-14
Mathematical Operators	C-15
Conditional Expressions	C-15
CASE (Switch)	C-15
CASE (If).....	C-16
Expressing Literals.....	C-17
Character Literals	C-17
Datetime Literals	C-18
Numeric Literals	C-18
Integer Literals	C-18
Decimal Literals	C-18
Floating Point Literals	C-19
Calculated Members	C-19
CALCULATEDMEMBER Syntax	C-19
Rules for the CALCULATEDMEMBER Expression.....	C-20
Using Solve Order to Control Formula Evaluation Sequence.....	C-21
Examples of Calculated Members in Queries.....	C-22
Variables	C-23
Aggregate, Running Aggregate, and Time Series Functions	C-24
Aggregate Functions.....	C-24
AGGREGATE AT.....	C-25
AVG	C-25
AVGDISTINCT	C-26
BOTTOMN.....	C-26
COUNT.....	C-26

COUNTDISTINCT.....	C-26
COUNT(*)	C-26
FIRST.....	C-27
GROUPBYCOLUMN	C-27
GROUPBYLEVEL	C-28
LAST	C-28
MAX.....	C-28
MEDIAN	C-29
MIN.....	C-29
NTILE	C-29
PERCENTILE.....	C-29
RANK	C-30
STDDEV	C-30
STDDEV_POP	C-30
SUM	C-31
SUMDISTINCT.....	C-31
TOPN	C-31
Running Aggregate Functions	C-31
MAVG.....	C-32
MSUM.....	C-32
RSUM.....	C-33
RCOUNT	C-34
RMAX	C-34
RMIN	C-35
Time Series Functions	C-35
AGO	C-36
Determining the Level Used by the AGO Function	C-36
PERIODROLLING.....	C-37
Determining the Level Used by the PERIODROLLING Function	C-38
TODATE.....	C-38
String Functions.....	C-39
ASCII.....	C-39
BIT_LENGTH	C-40
CHAR.....	C-40
CHAR_LENGTH.....	C-40
CONCAT	C-40
INSERT	C-41
LEFT	C-41
LENGTH.....	C-42
LOCATE	C-42
LOCATEN	C-43
LOWER.....	C-43
OCTET_LENGTH	C-43
POSITION	C-44
REPEAT	C-44
REPLACE	C-44
RIGHT	C-45

SPACE.....	C-45
SUBSTRING.....	C-46
TRIMBOTH.....	C-46
TRIMLEADING.....	C-46
TRIMTRAILING.....	C-46
UPPER.....	C-47
Math Functions	C-47
ABS.....	C-48
ACOS.....	C-48
ASIN.....	C-48
ATAN.....	C-48
ATAN2.....	C-48
CEILING.....	C-49
COS.....	C-49
COT.....	C-49
DEGREES.....	C-49
EXP.....	C-49
EXTRACTBIT.....	C-50
FLOOR.....	C-50
LOG.....	C-50
LOG10.....	C-50
MOD.....	C-51
PI.....	C-51
POWER.....	C-51
RADIANs.....	C-51
RAND.....	C-51
RANDFROMSEED.....	C-52
ROUND.....	C-52
SIGN.....	C-52
SIN.....	C-52
SQRT.....	C-53
TAN.....	C-53
TRUNCATE.....	C-53
Calendar Date/Time Functions	C-53
CURRENT_DATE.....	C-54
CURRENT_TIME.....	C-54
CURRENT_TIMESTAMP.....	C-54
DAY_OF_QUARTER.....	C-55
DAYNAME.....	C-55
DAYOFMONTH.....	C-55
DAYOFWEEK.....	C-55
DAYOFYEAR.....	C-56
HOUR.....	C-56
MINUTE.....	C-56
MONTH.....	C-56
MONTH_OF_QUARTER.....	C-56
MONTHNAME.....	C-57

NOW	C-57
QUARTER_OF_YEAR.....	C-57
SECOND.....	C-57
TIMESTAMPADD	C-57
TIMESTAMPDIFF	C-59
WEEK_OF_QUARTER.....	C-60
WEEK_OF_YEAR.....	C-60
YEAR.....	C-61
Conversion Functions	C-61
CAST	C-61
CHOOSE.....	C-62
IFNULL.....	C-62
INDEXCOL	C-62
Example With Hierarchy Levels.....	C-63
TO_DATETIME	C-64
VALUEOF	C-65
Database Functions	C-65
EVALUATE.....	C-66
EVALUATE_ANALYTIC	C-67
EVALUATE_AGGR.....	C-67
EVALUATE_PREDICATE.....	C-68
Hierarchy Navigation Functions	C-69
ISANCESTOR	C-69
ISCHILD	C-70
ISDESCENDANT	C-71
ISLEAF	C-72
ISPARENT	C-73
ISROOT	C-74
System Functions	C-75
USER	C-75
DATABASE.....	C-76

D Merge Rules

General Merge Rules and Behavior	D-1
Special Merge Algorithms for Logical Table Sources and Other Objects	D-4
Merging Objects that Use the Vector Merge Algorithm.....	D-4
Merging Logical Table Sources	D-6
Merging Security Filters	D-6
Inferring the Use Logical Column Property for Presentation Columns	D-6
Merging Aliases.....	D-7

E Deleting Unwanted Objects from the Repository

About the Object Pruning Utility	E-1
Using the Object Pruning Utility	E-1
Creating the Input File.....	E-1
Running the Prune Utility.....	E-2
Log File	E-3

Error Log File.....	E-3
Deletion Rules for the Object Pruning Utility	E-3

F Exchanging Metadata with Databases to Enhance Query Performance

About Exchanging Metadata with Databases	F-1
Generating the Import File	F-1
Running the Generator.....	F-2
About the Metadata Input File.....	F-4
About the Output Files.....	F-5
Troubleshooting Errors from the Generator	F-5
Metadata Conversion Rules and Error Messages.....	F-6
Conversion Rules for Oracle Databases.....	F-6
Conversion Rules for IBM DB2 Databases	F-7
Using Materialized Views in the Oracle Database with Oracle Business Intelligence	F-10
About Using the SQL Access Advisor with Materialized Views.....	F-10
Deploying Metadata for Oracle Database	F-11
Executing the SQL File for Oracle Database	F-11
Defining Constraints for the Existence of Joins	F-11
Creating the Query Workload.....	F-12
Creating Materialized Views.....	F-13
Using IBM DB2 Cube Views with Oracle Business Intelligence	F-14
About Using IBM DB2 Cube Views with Oracle Business Intelligence.....	F-14
Deploying Cube Metadata.....	F-15
Executing the Alias-SQL File for IBM Cube Views.....	F-15
Importing the XML File	F-15
Guidelines for Importing the XML File Using the IBM OLAP Center	F-15
Guidelines for Changing Cube Metadata After Importing the XML File	F-16
Guidelines for Creating Materialized Query Tables (MQTs).....	F-16

G Administration Tool Keyboard Shortcuts

Menu Keyboard Shortcuts	G-1
Dialog Keyboard Shortcuts	G-2
Physical Diagram and Business Model Diagram Keyboard Shortcuts	G-3

Glossary

Index

Preface

The Oracle Business Intelligence Foundation Suite is a complete, open, and integrated solution for all enterprise business intelligence needs, including reporting, ad hoc queries, OLAP, dashboards, scorecards, and what-if analysis. The Oracle Business Intelligence Foundation Suite includes Oracle Business Intelligence Enterprise Edition.

Oracle Business Intelligence Enterprise Edition (Oracle BI EE) is a comprehensive set of enterprise business intelligence tools and infrastructure, including a scalable and efficient query and analysis server, an ad-hoc query and analysis tool, interactive dashboards, proactive intelligence and alerts, and an enterprise reporting engine.

The components of Oracle BI EE share a common service-oriented architecture, data access services, analytic and calculation infrastructure, metadata management services, semantic business model, security model and user preferences, and administration tools. Oracle BI EE provides scalability and performance with data-source specific optimized request generation, optimized data access, advanced calculation, intelligent caching services, and clustering.

This guide contains information about building an Oracle Business Intelligence metadata repository and includes topics on setting up and connecting to data sources, building the Physical layer, Business Model and Mapping layer, and Presentation layer, how to use the multiuser development environment, and a Logical SQL reference.

Audience

This document is intended for anyone who intends to design and build a metadata repository using the Oracle Business Intelligence Administration Tool, such as a Business Intelligence strategist, metadata provider, or ETL developer.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/us/corporate/accessibility/index.html>.

Access to Oracle Support

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

Related Documents

For more information, see the following documents in the Oracle Business Intelligence Enterprise Edition 11g Release 1 (11.1.1) documentation set:

- The Oracle Business Intelligence chapter in *Oracle Fusion Middleware Release Notes* for your platform
- *Oracle Fusion Middleware Installation Guide for Oracle Business Intelligence*
- *Oracle Fusion Middleware Upgrade Guide for Oracle Business Intelligence*
- *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*
- *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*
- *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition*
- *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition*
- *Oracle Fusion Middleware Developer's Guide for Oracle Business Intelligence Enterprise Edition*
- *Oracle Fusion Middleware Scheduling Jobs Guide for Oracle Business Intelligence Enterprise Edition*

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.

New Features for Oracle BI Metadata Repository Builders

This preface describes changes to metadata repository features for Oracle Business Intelligence Enterprise Edition 11g Release 1 (11.1.1). If you are upgrading to Oracle BI EE from a previous release, read the following information carefully, because there are significant differences in features, tools, and procedures.

This preface contains the following topics:

- [New Features for Oracle BI EE 11g Release 1 \(11.1.1.5\)](#)
- [New Features for Oracle BI EE 11g Release 1 \(11.1.1.3\)](#)

New Features for Oracle BI EE 11g Release 1 (11.1.1.5)

This section contains the following topics:

- [New Features](#)
- [Upgrade Considerations](#)

New Features

New metadata repository features in Oracle BI EE 11g Release 1 (11.1.1.5) include:

- [Access to Oracle OLAP Data Sources](#)
- [Access to TimesTen Data Sources](#)
- [Ability to Connect to SAP/BW Data Sources Using a Native Connection](#)
- [New Diagramming Capability](#)
- [Validation Utility to Validate Metadata Objects in a Repository](#)

Access to Oracle OLAP Data Sources

Oracle BI EE now supports Oracle OLAP as a data source. See "[Working with Oracle OLAP Data Sources](#)" for more information.

Access to TimesTen Data Sources

Oracle BI EE now supports Oracle TimesTen In-Memory Database as a data source. See "[Setting Up Oracle TimesTen In-Memory Database Data Sources](#)" for more information.

Ability to Connect to SAP/BW Data Sources Using a Native Connection

You can now use the SAP BW Native connection option to connect to SAP/BW data sources over BAPI. See ["Setting Up SAP/BW Data Sources"](#) for more information.

New Diagramming Capability

The Physical and Business Model Diagrams have enhanced capabilities in this release, including an improved look-and-feel, the ability to display columns in table objects in the diagram, and new diagram options like auto-layout, marquee zoom, and zoom to fit. See ["Using the Physical and Business Model Diagrams"](#) for more information.

Validation Utility to Validate Metadata Objects in a Repository

Oracle BI EE now includes a validation utility to validate metadata objects. See ["Using the validaterpd Utility to Check Repository Consistency"](#) for more information.

Upgrade Considerations


Be aware of the following upgrade considerations when upgrading to Oracle BI EE 11g Release 1 (11.1.1.5):

- Joins in the Physical and Business Model Diagrams are represented by a line with an arrow at the "one" end of the join, rather than the line with crow's feet at the "many" end of the join that was used in previous releases.

For example, the following image shows a join in the diagram as it was represented in previous releases:



This join is represented as follows in 11g Release 1 (11.1.1.5):



- When creating joins in the Physical and Business Model Diagrams, you now select the "many" end of the join first, and then select the "one" end of the join. In previous releases, joins in the diagrams were created by selecting the "one" end of the join first.

This new gesture direction (from many to one) matches the direction of the new join arrow, described in the previous bullet.

- The Consistency Check Manager now provides a warning when the same connection pool is being used for both queries and for initialization blocks. This configuration is not recommended. Instead, create a dedicated connection pool for initialization blocks. Otherwise, query performance might suffer, or user logins might hang if authorization initialization blocks cannot run. These warnings appear similar to the following:

```
[39062] Initialization Block 'Authorization' uses Connection Pool 'My_DB'.  
"My_CP" which is used for report queries. This may impact query performance.
```

For more information about upgrading to Oracle BI EE 11g, see *Oracle Fusion Middleware Upgrade Guide for Oracle Business Intelligence*.

New Features for Oracle BI EE 11g Release 1 (11.1.1.3)

This section contains the following topics:

- [New Features](#)
- [Upgrade Considerations](#)

New Features

New metadata repository features in Oracle BI EE 11g Release 1 (11.1.1.3) include:

- [New SampleApp.rpd Demonstration Repository](#)
- [Hierarchy Objects in the Presentation Layer](#)
- [Support for Unbalanced \(Ragged\) and Skip-Level Hierarchies](#)
- [Support for Parent-Child Hierarchies](#)
- [Creating Subject Areas for Logical Stars and Snowflakes](#)
- [Multicurrency Support](#)
- [Generating and Applying XML Patch Files](#)
- [Enhancements for Multiuser Development and Repository Lifecycle Management](#)
- [Metadata Repository File Encryption](#)
- [Metadata Repository File Compression](#)
- [Embedded Database Functions](#)
- [SQL Functions for Time Series Calculations and Level-Based Measures](#)
- [Support for SELECT_PHYSICAL](#)
- [Support for Calculated Members](#)
- [Access to Essbase Data Sources](#)
- [Access to Hyperion Financial Management Data Sources](#)
- [Access to ADF Business Component Data Sources](#)
- [Access to Oracle RPAS Data Store Through ODBC](#)
- [Ability to Connect to Oracle Database Through OCI](#)
- [Support for Dragging and Dropping XMLA Objects](#)
- [Support for Using a Standby Database with Oracle Business Intelligence](#)
- [Aggregate Persistence Auto-Indexing](#)
- [Pruning Utility to Delete Unwanted Repository Objects](#)
- [Ability to Defer Execution of Session Variable Initialization Blocks](#)
- [Improved Usability in the Administration Tool](#)

New SampleApp.rpd Demonstration Repository

A new demonstration repository called SampleApp.rpd is available for this release of Oracle BI EE. This repository provides best practice information about modeling many different types of objects.

A basic version of SampleApp.rpd, called SampleAppLite.rpd, is automatically installed as the default repository when you choose the Simple Install option. The full version of SampleApp.rpd, which contains many additional examples and features, is

available on the Oracle Technology Network. See ["About the SampleApp.rpd Demonstration Repository"](#) for more information.

Hierarchy Objects in the Presentation Layer

You can now define presentation hierarchies and presentation levels in the Presentation layer. These objects provide an explicit way to expose the multidimensional model in Oracle BI Answers and enables users to create hierarchy-based queries. Presentation hierarchies expose analytic functionality such as member selection, custom member groups, and asymmetric queries. See ["Working with Presentation Hierarchies and Levels"](#) for more information.

Support for Unbalanced (Ragged) and Skip-Level Hierarchies

Oracle BI EE now supports unbalanced and skip-level hierarchies. An unbalanced (or ragged) hierarchy is a hierarchy where the leaves (members with no children) do not necessarily have the same depth. A skip-level hierarchy is a hierarchy where there are members that do not have a value for a particular ancestor level. See ["About Level-Based Hierarchies"](#) for more information.

Support for Parent-Child Hierarchies

Oracle BI EE now supports parent-child hierarchies. Parent-child hierarchies (also called value hierarchies) contain members that all have the same type. For example, an organizational chart has a distinct parent-child hierarchy, but all members are employees. See ["Creating and Managing Dimensions with Parent-Child Hierarchies"](#) for more information.

Creating Subject Areas for Logical Stars and Snowflakes

You can automatically create one subject area (formerly called presentation catalog) for each logical star or logical snowflake in your business model. See ["Automatically Creating Subject Areas Based on Logical Stars and Snowflakes"](#) for more information.

Multicurrency Support

You can configure logical columns so that Oracle BI EE users can select the currency in which they prefer to view currency columns in analyses and dashboards. See ["Configuring Logical Columns for Multicurrency Support"](#) for more information.

Generating and Applying XML Patch Files

You can now generate an XML patch file that contains only the changes made to a repository. This patch can be then applied to the old (original) version of the repository to create the new version. This is very useful for development-to-production scenarios, and can also be used for Oracle BI Applications customers to upgrade their repository. See ["Performing Patch Merges"](#) for more information.

You can also use the Oracle BI Server XML utilities to create a generic, XML-based representation of the Oracle BI repository metadata, on any supported Oracle BI Server operating system. See ["About the Oracle BI Server XML API"](#) in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Enhancements for Multiuser Development and Repository Lifecycle Management

The following list summarizes improvements and new features for multiuser development and repository lifecycle management:

- MUD projects now explicitly include subject areas (formerly called presentation catalogs). Object security permissions no longer affect which presentation objects are included in a checked-out repository.
- You can select options in the Merge Repository Wizard or set options in the MUD options file to automatically check consistency and equalize during merges.
- Improved Merge Repository Wizard integration enables a single step input for repositories. You no longer have to select repositories one by one and wait for them to load.
- You can now extract objects on all platforms using the command-line utility `extractprojects`.
- The two-way merge feature (full merge without a common parent) enables easier consolidation of separate repositories.

Metadata Repository File Encryption

Repositories now have a repository password that encrypts the repository contents using a strong encryption algorithm. This feature ensures that repository metadata is secure, including data source credentials in connection pool objects. See ["Changing the Repository Password"](#) for more information.

Metadata Repository File Compression

To reduce storage needs, repositories are now stored in a compressed format. Because of this, you may notice that the size of an RPD file opened and saved in this release is significantly smaller than the size of RPD files from previous releases. See ["About the Multiuser Development Environment"](#) for more information about using and storing multiple repositories.

Embedded Database Functions

Users and administrators can create requests by directly calling database functions from either Oracle BI Answers, or by using a Logical column (in the Logical Table source) within the metadata repository. These functions provide the ability to access custom written functions or procedures on the underlying database. Supported functions include `EVALUATE`, `EVALUATE_AGGR`, and `EVALUATE_PREDICATE`. See ["Database Functions"](#) for more information.

In addition, Evaluate functions can be leveraged significantly for Essbase users. See ["Examples Using EVALUATE_AGGREGATE and EVALUATE to Leverage Unique Essbase Functions"](#) for more information.

SQL Functions for Time Series Calculations and Level-Based Measures

The following new and revised SQL functions enable users to create time series calculations and level-based measures:

- `PERIODROLLING` is a new function that computes the aggregate of a measure over the period starting x units of time and ending y units of time from the current time.
- `AGGREGATE AT` is a new function that aggregates columns based on the level or levels specified.
- `AGO` is a time series aggregation function that calculates the aggregated value from the current time back to a specified time period. In this release, the `time_level` argument is optional, and there are additional syntax changes. This function was called `PERIODAGO` in a previous release.

- `TODATE` is a time series aggregation function that aggregates a measure attribute from the beginning of a specified time period to the currently displayed time. The syntax for this function has changed in this release. This function was called `PERIODTODATE` in a previous release.

See ["Aggregate Functions"](#) for more information about these functions. See also ["Modeling Time Series Data"](#) for information about using the time series functions `AGO`, `TODATE`, and `PERIODROLLING`.

Support for `SELECT_PHYSICAL`

`SELECT_PHYSICAL` statements provide the functionality to directly query objects in the Physical layer of the metadata repository, and to nest such a statement within a query against the Business Model and Mapping layer or the Presentation layer. See ["Syntax and Usage Notes for `SELECT_PHYSICAL`"](#) for more information.

Support for Calculated Members

Calculated members are user-defined dimension members whose measure values are calculated at run time. You can define a calculated member within a dimension through a formula that references other members of the same dimension. See ["Calculated Members"](#) for more information.

Access to Essbase Data Sources

Oracle BI EE now supports Essbase as a data source. See ["Working with Essbase Data Sources"](#) for more information.

Access to Hyperion Financial Management Data Sources

Oracle BI EE now supports Hyperion Financial Management as a data source. See ["Working with Hyperion Financial Management Data Sources"](#) for more information.

Access to ADF Business Component Data Sources

Oracle BI EE now supports using ADF Business Components as data sources. With this feature, users can integrate operational reporting with any application that is built on top of the ADF Framework. See ["Working with ADF Business Component Data Sources"](#) for more information.

Access to Oracle RPAS Data Store Through ODBC

The Oracle BI Server can now access Oracle Retail Predictive Application Server (RPAS) data stores using the RPAS ODBC driver. See ["Setting Up Oracle RPAS Data Sources"](#) for more information.

Ability to Connect to Oracle Database Through OCI

Administrators can now use the native Oracle Call Interface (OCI) when importing physical objects from an Oracle Database. Because using OCI is faster than using an ODBC connection, administrators should always use OCI when importing from or setting up connections to an Oracle Database. See ["Setting Up Oracle Database Data Sources"](#) for more information.

Support for Dragging and Dropping XMLA Objects

You can now drag and drop XMLA objects (whole XMLA database or cube tables, at the same time or one by one) into the logical layer to automatically create a consistent business model. See ["Creating the Business Model and Mapping Layer"](#) and ["Automatically Creating Business Model Objects for Multidimensional Data Sources"](#) for more information.

Support for Using a Standby Database with Oracle Business Intelligence

You can now use a standby database with Oracle BI EE. A standby database is used mainly for its high availability and failover functions as a backup for the primary database. See ["Using a Standby Database with Oracle Business Intelligence"](#) for more information.

Aggregate Persistence Auto-Indexing

For greater performance, Oracle BI EE now creates indexes and statistics on relational tables when aggregates are persisted. See [Chapter 12](#) for more information about aggregate persistence.

Pruning Utility to Delete Unwanted Repository Objects

Oracle BI EE now includes a pruning utility to remove unwanted objects from your repository. This utility is especially useful for Oracle BI Applications customers who only need to use a subset of the objects included in the shipped Oracle BI Applications repository. See [Appendix E, "Deleting Unwanted Objects from the Repository"](#) for more information.

Ability to Defer Execution of Session Variable Initialization Blocks

To improve session logon time and save resources, administrators can defer execution of selected session variable initialization blocks until the associated session variables are actually used within the session. See ["Creating Initialization Blocks"](#) for more information.

Improved Usability in the Administration Tool

Many Administration Tool screens, wizards, and other parts of the user interface were updated and streamlined for this release of Oracle BI EE. Updated wizards include the Create New Repository Wizard, Aggregate Persistence Wizard, and Merge Wizard. Expression Builder and the Consistency Checker also contain usability enhancements in this release.

Additional changes include updated and improved icons, a simpler process flow to create joins, improvements to the Query Repository feature, and a simplified way to add some repository objects. Also, shortcuts were added to enable keyboard access to menu options, dialog navigation, and navigation in the Physical and Business Model Diagrams.

Upgrade Considerations

Be aware of the following upgrade considerations when upgrading to Oracle BI EE 11g Release 1 (11.1.1.3):

- You must upgrade your repository files from previous releases before they will work in this release of Oracle BI EE. See *Oracle Fusion Middleware Upgrade Guide for Oracle Business Intelligence* for full information about upgrading a repository from a previous release.
- Many configuration settings that affect repository development, including the default published repository, are now centrally managed in Fusion Middleware Control. You can no longer manually change these configuration settings in NQSCONFIG.INI. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.
- Be aware of the following security-related changes:
 - Repositories now have repository-specific passwords that are used to encrypt the repository contents. The repository password is stored in an external

credential store when you publish a repository in Fusion Middleware Control, so that the Oracle BI Server can retrieve the password to load the repository. See ["Changing the Repository Password"](#) for more information.

Note that a blank repository password is not allowed.

- Groups no longer exist in the repository as objects. Instead, you implement data access security based on the application roles to which a user belongs.
Application roles are managed in an external policy store. Application role objects exist in the repository, but these objects are pointers (references) to the externally managed roles.
- Users are managed in an external authentication provider and are no longer managed in the repository. User objects exist in the repository, but these objects are pointers (references) to the externally managed users.

See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for information about these and other security changes. See also [Chapter 13, "Applying Data Access Security to Repository Objects"](#) for additional information.

- The connection parameters in the default ODBC DSN for the Oracle BI Server are now centrally managed by Fusion Middleware Control and cannot be manually changed.

In addition, Oracle BI EE is now deployed in a clustered configuration by default. Because of this, the default ODBC DSN for the Oracle BI Server points to the Cluster Controller by default, rather than to the Oracle BI Server.

See ["Integrating Other Clients with Oracle Business Intelligence"](#) in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about ODBC DSNs for the Oracle BI Server.

- This release of Oracle BI EE has additional dependencies for a running system, including:
 - The relational database specified upon installation must be running. This database must contain required Oracle BI EE schemas loaded using the Repository Creation Utility (RCU).
 - Oracle WebLogic Server must be running.
- You can no longer restart the Oracle BI Server using the Administration Tool in online mode. Instead, you can restart the Oracle BI Server and other system processes using Fusion Middleware Control and the Oracle WebLogic Server Administration Console.

You can also use the BI Systems Management API to programmatically start and stop Oracle BI EE. This feature is especially helpful for automating rolling restart of Oracle BI Servers in a cluster, to enable repository upgrade with zero end-user downtime.

For more information, see ["Starting and Stopping Oracle Business Intelligence"](#) and ["Starting and Stopping Oracle Business Intelligence Using the Oracle BI Systems Management API"](#) in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

- Before you can run any of the Oracle BI Server command-line utilities, you must now run `bi-init.cmd` (or `bi-init.sh` on UNIX) to launch a command prompt or shell window that is initialized to your Oracle instance. See ["Running bi-init to Launch a Shell Window Initialized to Your Oracle Instance"](#) for more information.

- If you want to use a net service name in a connection pool for an Oracle Database data source, you must set up a tnsnames.ora file in the following location within the Oracle BI EE environment, so that the Oracle BI Server can locate the entry:
ORACLE_HOME/network/admin
- You can no longer open the Administration Tool by double-clicking a repository file. The resulting Administration Tool window is not initialized to your Oracle instance, and errors will result later in your session. Instead, always use the Start menu to open the Administration Tool, or launch the Administration Tool from the command line using bi-init.cmd. See "[Opening the Administration Tool](#)" for more information.
- Presentation catalogs in the Presentation layer are now called subject areas.
- Static repository variables must now have default initializers that are constant values. See "[About Repository Variables](#)" for more information.
- If you configured write-back capability in a previous release of Oracle Business Intelligence, you must now explicitly select the **Writeable** option for each logical column for which you want to enable write-back. See "[Enabling Write Back On Columns](#)" for more information.
- Bridge tables are now identified using repository modeling techniques. They are no longer identified using the **Bridge table** option in the Logical Table dialog that existed in previous releases. Check your repository to ensure that your bridge tables are modeled appropriately. See "[Modeling Bridge Tables](#)" for more information.
- You might notice that some queries that used to return one result now return a different result. This behavior occurs because the determination of which logical table source to use for a query is now very ordered, whereas in previous releases, the determination was random. Examine and adjust the modeling to correct the behavior.
- In the previous release, if two logical table sources for the same logical table map to the same physical table, and both logical table sources are used in a query, and both logical table sources have a WHERE clause filter, then the filter from only one of the logical table sources was applied. The other WHERE clause filter was ignored.

In the current release, in this situation, the WHERE clause filters from both logical table sources are applied to the query. Typically, this behavior produces the desired results. If you encounter errors related to this issue, you can correct them by using physical table aliases to ensure that the same physical table is not mapped to the same logical table at different levels.
- The Consistency Check Manager now enforces additional validation rules to help ensure that your repository is consistent. In addition, some rules that existed in previous releases might now be displayed during consistency checks. The following table summarizes these rules:

Validation Rule Example	Type	Description
[14031] The content filter of a source for logical table: FACT_TABLE_NAME references multiple dimensions.	Error	The given logical table has a logical table source with a WHERE clause filter that references multiple dimensions. A WHERE clause with multiple dimensions is invalid.

Validation Rule Example	Type	Description
[38126] 'Logical Table' ""Technology - WFA"."Fact WFA WO "" has name with leading or trailing space(s).	Error	Identifies an object with leading or trailing spaces in the object name. Repository objects can no longer have leading or trailing spaces in their names. Leading and trailing spaces in object names can cause query and reporting issues.
[38012] Logical column DIM_Start_Date.YEAR_QUARTER_NBR does not have a physical data type mapping, nor is it a derived column. [38001] Logical column DIM_Start_Date.YEAR_QUARTER_NBR has no physical data source mapping.	Error	Logical columns that are mapped to logical table sources that are disabled are reported as consistency errors, because the logical table source mappings are invalid and would cause queries to fail. Both of the given validation rules are related to the same issue.
[39028] The features in Database 'MyDB' do not match the defaults. This can cause query problems.	Warning	Some database feature defaults were changed in this release of Oracle BI EE. Unless you have specific customizations to your feature set, it is recommended that you reset your database features to the new defaults.
[39003] Missing functional dependency association for column: DIM_Offer_End_Date.CREATE_DT.	Warning	This warning indicates that the given column is not associated to any level. When this situation occurs, the column will be associated by default to the lowest level in the parent dimension. The warning brings this issue to the repository developer's attention in case the default behavior is not desired.
[39009] Logical fact table MY_FACT has an aggregate source MY_FACT.YRLY_SUM that does not join to a Dimension: Hier - Time source at the proper level of detail. [39055] Fact table "HR"."FACT - HC Budget" is not joined to tables in logical dimension "HR"."DIM - HR EmployeeDim". This will cause problems when extracting project(s). [39059] Logical dimension table MY_DIM has a source MY_DIM_DAILY at level Daily that joins to a higher level fact source MY_FACT_SUM.MTHLY_SUM	Warning	Even though this fact logical table source has an aggregate grain set in this dimension, no valid physical join was found that connects to any logical table source in this dimension. This means that either no join exists at all, or it does exist but is invalid because it connects a higher-level fact source to a lower-level dimensional source. Such joins are invalid and ignored by the Oracle BI Server because, if followed, they would lead to double counting in query answers. For example, consider Select year, yearlySales. Even if a join exists between monthTable and yearlySales table on yearId, it cannot be used because such a join would overstate the results by a factor of 12 (the number of months in each year). All three of the given validation rules are related to the same issue.
[39054] Fact table "Sales - STAR"."Fact - STAR Statistics" is not joined to logical dimension table "Sales - STAR"."Dim - Plan". This will cause problems when extracting project(s).	Warning	This warning indicates that the aggregation content filter "Group by Level" in the logical table source of a fact table references logical dimension tables that are not joined to that fact table. If that fact table is extracted in the extract/MUD process, the dimensions that are not joined will not be extracted. In this case, the aggregation content of the extracted logical table source would not be the same as in the original logical table source.
[39057] There are physical tables mapped in Logical Table Source ""HR"."Dim - Schedule"."SCH_DEFN"" that are not used in any column mappings or expressions.	Warning	This warning indicates that the given logical table source has irrelevant tables added that are not used in any mapping. This situation will not cause any errors.

- Invalid objects are now deleted during Consistency Checks. This behavior might result in deleted expressions and filters on logical table sources and logical columns. Invalid references can occur when objects were deleted in the Physical

layer without properly accounting for the references in the Business Model and Mapping layer objects.

- See "[About Converting Older Projects During Repository Upgrade](#)" for information about how projects are upgraded when you upgrade a repository from Oracle BI EE versions before 10.1.3.2.
- The environment variable `OBIS_Essbase_CustomGroup_Generation`, used in previous releases to customize the use of custom group syntax with Essbase, has been replaced by a new database feature called `PERF_CUSTOM_GROUP_GENERATION_MODE`. This database feature impacts how custom group syntax is generated on Essbase and other multidimensional sources. The set of valid values is the same as for the environment variable (0-2).
- The environment variable `OBIS_Essbase_NonEmptyTuples_Generation.Database.Catalog.CubeTable`, used in previous releases to resolve issues with large query sets, has been replaced by a new database feature called `PERF_PREFER_SUPPRESS_EMPTY_TUPLES`. This database feature controls whether empty tuples with empty cell values are eliminated. Note that this database feature does not change the null suppression behavior on the final result set.

For more information about upgrading to Oracle BI EE 11g, see *Oracle Fusion Middleware Upgrade Guide for Oracle Business Intelligence*.

Introduction to Building Your Metadata Repository

This chapter explains how to plan and design your metadata repository, including how to plan your business model, how to work with the physical content for your business model, and general repository design guidelines.

To effectively plan and build your metadata repository, you need to have experience with SQL queries and be familiar with reporting and analysis. You should also have experience with industry-standard data warehouse modeling practices, and be familiar with general relational entity-relationship modeling.

This chapter contains the following topics:

- [About Oracle BI Server and Oracle BI Repository Architecture](#)
- [Planning Your Business Model](#)
- [Identifying the Data Source Content for the Physical Layer](#)
- [Guidelines for Designing a Repository](#)
- [Topics of Interest in Other Guides](#)
- [System Requirements and Certification](#)

About Oracle BI Server and Oracle BI Repository Architecture

The architecture of the Oracle BI Server and the Oracle BI repository provides a layer of abstraction that lets users send simple Logical SQL queries against complex federated data sources.

This section contains the following topics:

- [About Oracle BI Server Architecture](#)
- [About Layers in the Oracle BI Repository](#)

About Oracle BI Server Architecture

The Oracle BI Server is an Oracle Business Intelligence component that processes user requests and queries underlying data sources. The Oracle BI Server maintains the logical data model and provides client access to this model through ODBC.

The Oracle BI Server uses the metadata in the Oracle BI repository to perform the following two tasks:

- Interpret Logical SQL queries and write corresponding physical queries against the appropriate data sources


- Transform and combine the physical result sets and perform final calculations

The Oracle BI Server connects to the underlying data sources through either ODBC or through native APIs, such as OCI for Oracle Database.

The Administration Tool client is a Windows application that you can use to create and edit your Oracle BI repository. The Administration Tool can connect directly to the repository in offline mode, or it can connect to the repository through the Oracle BI Server. Some options are only available in online mode. See ["Using Online and Offline Repository Modes"](#) for more information.

Figure 1–1 shows how the Oracle BI Server interacts with query clients, data sources, the Oracle BI repository, and the Administration Tool.

Figure 1–1 Oracle BI Server Architecture



Example 1–1 shows how the Oracle BI Server interprets and converts Logical SQL queries.

Example 1–1 Logical Requests Are Transformed Into Complex Physical Queries

Assume the Oracle BI Server receives the following simple client request:

```
SELECT
"D0 Time"."T02 Per Name Month" saw_0,
"D4 Product"."P01 Product" saw_1,
"F2 Units"."2-01 Billed Qty (Sum All)" saw_2
FROM "Sample Sales"
ORDER BY saw_0, saw_1
```

The Oracle BI Server can then convert the Logical SQL query into a sophisticated physical query, as follows:

```
WITH SAWITH0 AS (
select T986.Per_Name_Month as c1, T879.Prod_Dsc as c2,
sum(T835.Units) as c3, T879.Prod_Key as c4
from
Product T879 /* A05 Product */ ,
Time_Mth T986 /* A08 Time Mth */ ,
FactsRev T835 /* A11 Revenue (Billed Time Join) */
where ( T835.Prod_Key = T879.Prod_Key and T835.Bill_Mth = T986.Row_Wid)
```

```

group by T879.Prod_Dsc, T879.Prod_Key, T986.Per_Name_Month
)
select SAWITH0.c1 as c1, SAWITH0.c2 as c2, SAWITH0.c3 as c3
from SAWITH0
order by c1, c2

```

About Layers in the Oracle BI Repository

An Oracle BI repository has the following layers:

- **Physical layer.** This layer defines the objects and relationships that the Oracle BI Server needs to write native queries against each physical data source. You create this layer by importing tables, cubes, and flat files from your data sources.

Separating the logical behavior of the application from the physical model provides the ability to federate multiple physical sources to the same logical object, enabling aggregate navigation and partitioning, as well as dimension conformance and isolation from changes in the physical sources. This separation also enables the creation of portable BI Applications.

- **Business Model and Mapping layer.** This layer defines the business or logical model of the data and specifies the mapping between the business model and the physical schemas. This layer determines the analytic behavior seen by users, and defines the superset of objects and relationships available to users. It also hides the complexity of the source data models.

Each column in the business model maps to one or more columns in the Physical layer. At run time, the Oracle BI Server evaluates Logical SQL requests against the business model, and then uses the mappings to determine the best set of physical tables, files, and cubes for generating the necessary physical queries. The mappings often contain calculations and transformations, and might combine multiple physical tables.

- **Presentation layer.** This layer provides a way to present customized, secure, role-based views of a business model to users. It adds a level of abstraction over the Business Model and Mapping layer and provides the view of the data seen by users building requests in Presentation Services and other clients.

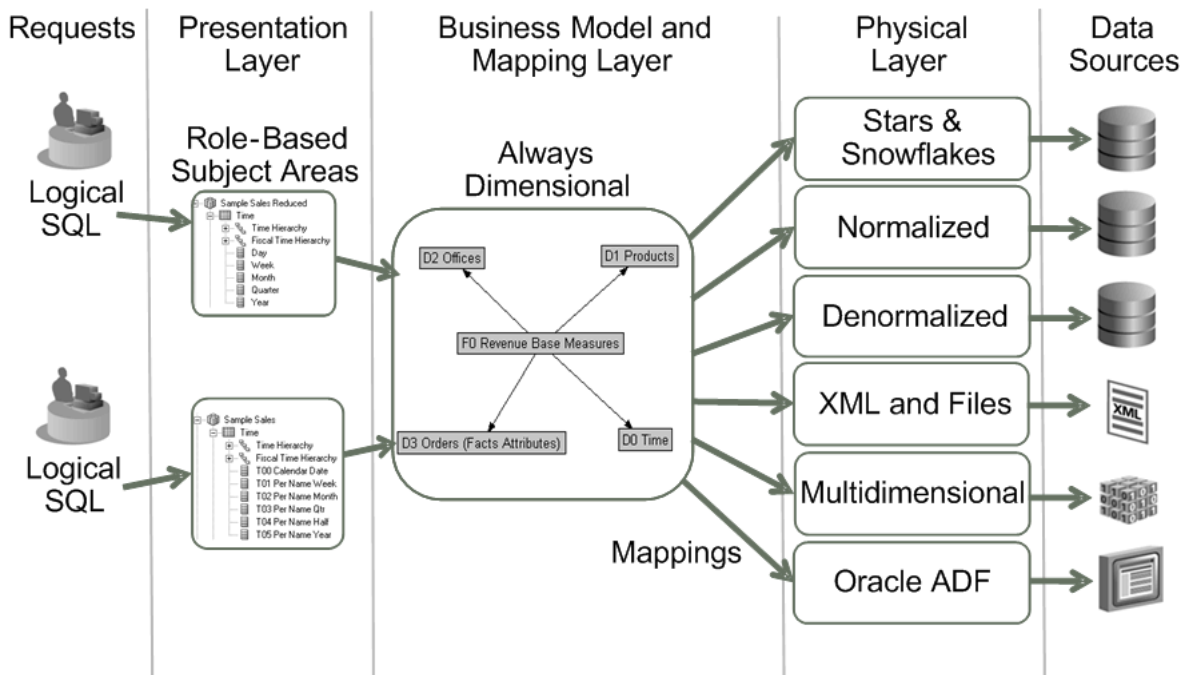
You can create multiple subject areas in the Presentation layer that map to a single business model, effectively breaking up the business model into manageable pieces.

Before you build any repository layers in the Administration Tool, it is important to create a high-level design of the Business Model and Mapping layer based on the analytic requirements of your users. After you have a conceptual design to work toward, you can then build your metadata objects.

The typical order is to create the Physical layer objects first, the Business Model and Mapping layer objects next, and the Presentation layer objects last. However, you can work on each layer at any stage. After you complete all three layers, you can set up security when you are ready to begin testing the repository.

Figure 1–2 shows how a Logical SQL query traverses the layers of an Oracle BI repository.

Figure 1–2 Logical SQL Query Traversing the Layers in an Oracle BI Repository



Note that a single Oracle BI repository can contain two or more independent semantic models, rather than a single, integrated, enterprise-wide model. A semantic model consists of one business model, its related objects in the Presentation and Physical layers, and additional related objects like variables, initialization blocks, and application roles. A semantic model is also known as a Common Enterprise Information Model.

See also [Figure A–4](#) for a visual representation of multiple semantic models.

Planning Your Business Model

Planning your business model is the first step in developing a usable data model for decision support. After you have followed the planning guidelines in this section, you can begin to create your repository.

Analyzing Your Business Model Requirements

Your first task is to thoroughly understand your business model requirements. You must first understand what business model you want to build before you can determine what the Physical layer needs to have in it.

In a decision support environment, the objective of data modeling is to design a model that presents business information in a manner that parallels business analysts' understanding of the business structure. A successful model allows the query process to become a natural process by enabling analysts to structure queries in the same intuitive fashion as they would ask business questions. This model must be one that business analysts inherently understand and that answers meaningful questions correctly.

Unlike visual SQL tools such as Oracle BI Publisher, the business model defines the analytic behavior of your BI application. In contrast, the Physical layer only provides the components used to assemble a physical query mapped from business model logic.

This requires breaking down your business into several components to answer the following questions:

- What kinds of business questions are analysts trying to answer?
- What are the measures required to understand business performance?
- What are all the dimensions under which the business operates? Or, in other words, what are the dimensions used to break down the measurements and provide headers for the reports?
- Are there hierarchical elements in each dimension, and what types of relationships define each hierarchy?

After you have answered these questions, you can identify and define the elements of your business model.

Identifying the Content of the Business Model

To determine what content to include in your business model, you must first identify the logical columns on which users need to query. Then, to establish the role played by each column, identify whether it is a measure column or a dimensional attribute. Finally, arrange the logical columns in a dimensional model based on the relevant roles, relationships between columns, and logic.

Businesses are analyzed by relevant dimensional criteria, and the business model is developed from these relevant dimensions. These dimensional models form the basis of the valid business models to use with the Oracle BI Server.

Although not all dimensional models are built around a star schema, it is a best practice to use a simple star schema in the business model layer. In other words, the dimensional model should represent some measurable facts that are viewed in terms of various dimensional attributes.

After you analyze your business model requirements, you need to identify the specific logical tables and hierarchies that you need to include in your business model.

This section contains the following topics:

- [Identifying Logical Fact Tables](#)
- [Identifying Logical Dimension Tables](#)
- [Identifying Dimensions](#)
- [Identifying Lookup Tables](#)

Identifying Logical Fact Tables

Logical fact tables in the Business Model and Mapping layer contain measures that have aggregations built into their definitions. Logical fact tables are different from physical fact tables in relational models, which instead define facts at the lowest grain of the table.


Measures aggregated from facts must be defined in a logical fact table. Measures are typically calculated data such as dollar value or quantity sold, and they can be specified in terms of dimensions. For example, you might want to determine the sum of dollars for a given product in a given market over a given time period.

Each measure has its own aggregation rule such as *SUM*, *AVG*, *MIN*, or *MAX*. A business might want to compare values of a measure and need a calculation to express the comparison. Also, aggregation rules can be specific to particular dimensions. The Oracle BI Server lets you define complex, dimension-specific aggregation rules.

You do not explicitly label tables in the Business Model and Mapping layer as fact tables or dimension tables. Rather, the Oracle BI Server treats tables at the "one" end of a join as dimension tables, and tables at the "many" end of a join as fact tables.

Figure 1–3 illustrates the many-to-one joins to a fact table in a Business Model Diagram. In the diagram, all joins have an arrow (indicating the one side) pointing away from the Fact-Pipeline table; no joins are pointing to it. For an example of this in a business model, open a repository in the Administration Tool, right-click a fact table, and select **Business Model Diagram > Whole Diagram**.

Figure 1–3 Diagram of Fact Table Joins



Identifying Logical Dimension Tables

A business uses facts to measure performance by well-established dimensions, for example, by time, product, and market. Every dimension has a set of descriptive attributes. Dimension tables contain attributes that describe business entities (such as Customer Name, Region, Address, Country and so on). Dimension tables also contain primary keys that identify each member. Unlike logical fact tables, which are different from physical fact tables in relational models, logical dimension tables behave very much like relational dimension tables.

Dimension table attributes provide context to numeric data, such as being able to categorize Service Requests. Attributes stored in this dimension might include Service Request Owner, Area, Account, Priority, and so on.

Dimensions in the business model are conformed dimensions. In other words, even if a particular data source has five different instances of a particular Customer table, the business model should only have one table. To achieve this, all five physical source instances of Customer are mapped to a single Customer logical table, with transformations in the logical table source as necessary. Conformed dimensions hide the complexity of the Physical layer from users and enable you to combine data from multiple fact sources at different grains. They also enable you to federate multiple data sources.

Also note that dimension and level keys in the business model should be business keys rather than generated surrogate keys. In other words, use "Customer Name" with values like "Oracle" instead of "Customer Key" with values like "1076823." Using business keys in the business model ensures that all sources for that dimension can be

conformed to the same logical dimension table with the same logical key and level key.

Although generated surrogate keys can still exist in the Physical layer, where they are useful for their query performance advantages on joins, you typically do not have surrogate key columns in the Business Model and Mapping layer at all.

Identifying Dimensions

Dimensions are categories of attributes by which the business is defined. Common dimensions are time periods, products, markets, customers, suppliers, promotion conditions, raw materials, manufacturing plants, transportation methods, media types, and time of day. Within a given dimension, there may be many attributes. For example, the time period dimension can contain the attributes day, week, month, quarter, and year. Exactly what attributes a dimension contains depends on the way the business is analyzed.

Dimensions typically contain hierarchies, which are sets top-down relationships between members within a dimension. There are two types of hierarchies: level-based hierarchies (structure hierarchies), and parent-child hierarchies (value hierarchies). Level-based hierarchies are those in which members of the same type occur only at a single level, while members in parent-child hierarchies all have the same type. Oracle Business Intelligence also supports a special type of level-based dimension, called a time dimension, that provides special functionality for modeling time series data.

In level-based hierarchies, levels roll up from lower level to higher level; for example, months can roll up into a year. These rollups occur over the hierarchy elements and span natural business relationships.

In parent-child hierarchies, the business relationships occur between different members of the same real-world type, such as the manager-employee relationship in an organizational hierarchy tree. Parent-child hierarchies do not have explicitly named levels. There is no limit to the number of implicit levels in a parent-child hierarchy.

To define your hierarchies, you define the "contains" relationships in your business (geographical, product, time, and so on) to drive rollup aggregations in all calculations, as well as drill-down navigation in reports and dashboards. For example, if month rolls up into year and an aggregate table exists at the month level, that table can be used to answer questions at the year level by adding up all of the month-level data for a year.

It is important to use the right type of hierarchy for your needs. To determine which type to use, consider the following:

- Are all the members of the same type (such as employee, assembly, or account), or are they of different types that naturally fall into levels (such as year-quarter-month, continent-country-state/province, or brand-line-product)?
- Do members have the same set of attributes? For example, in a parent-child hierarchy like Employees, all members might have a Hire Date attribute. In a level-based hierarchy like Time, however, the Day type might have a Holiday attribute, while Month does not.
- Are the levels fixed at design time (year-quarter-month), or can run-time business transactions add or subtract levels? For example, a level could be added when the current lowest-level employee hires a subordinate, who then becomes the new lowest level.
- Are there constraints in your primary data source that require a certain hierarchy type? If your primary data source is modeled in one way or the other, you might

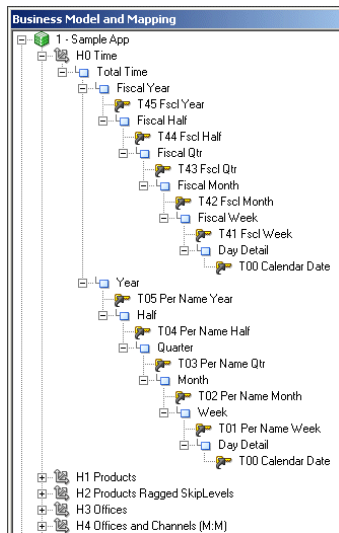
need to use the same hierarchy type in your business model, regardless of other factors.

See [Chapter 9, "Working with Logical Dimensions"](#) for more information.

About Dimensions with Multiple Hierarchies Sometimes, dimensions can contain multiple hierarchies. For example, time dimensions often have one hierarchy for the calendar year, and another hierarchy for the fiscal year. Note that dimensions with multiple hierarchies must always end with the same leaf table.

[Figure 1–4](#) shows a dimension with multiple hierarchies in the Business Model and Mapping layer of the Administration Tool.

Figure 1–4 Dimension with Multiple Hierarchies



Identifying Lookup Tables

When you need to display translated field information from multilingual schemas, you create a logical lookup table that corresponds to a lookup table in the Physical layer. A lookup table stores multilingual data corresponding to rows in the base tables. Before you can use a particular logical lookup table, you must designate it as a lookup table in the General tab of the Logical Table dialog. See "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization and lookup tables.

In addition to localization, lookup tables can be used any time you need to display one set of values to users, while using a different, corresponding set of values in the physical query. If necessary, the human-readable value can be looked up in a completely different data source.

Identifying the Data Source Content for the Physical Layer

After you have determined the requirements for your business model, you can look at what data source content you need in the Physical layer. Unlike the Business Model and Mapping layer, which is always dimensional, each physical model mirrors the shape of the source (for example, normalized, cube, and so on).

This section contains the following topics:

- [About Types of Physical Schemas in Relational Data Sources](#)
- [About Cubes in Multidimensional Data Sources](#)
- [Identifying the Data Source Table Structure](#)

About Types of Physical Schemas in Relational Data Sources

You can successfully model any physical schema in the Oracle BI repository, regardless of its type, because the model of any physical source can be broken down into overlapping subsets that are dimensional.

There are four types of physical schemas (models):

- **Star Schemas.** A star schema is a set of dimensional schemas (stars) that each have a single fact table with foreign key join relationships to several dimension tables. When you map a star to the business model, you first map the physical fact columns to one or more logical fact tables. Then, for each physical dimension table that joins to the physical fact table for that star, you map the physical dimension columns to the appropriate conformed logical dimension tables.
- **Snowflake Schemas.** A snowflake schema is similar to a star schema, except that each dimension is made up of multiple tables joined together. Like star schemas, you first map the physical fact columns to one or more logical tables. Then, for each dimension, you map the snowflaked physical dimension tables to a single logical table. You can achieve this by either having multiple logical table sources, or by using a single logical table source with joins.
- **Normalized Schemas.** Normalized schemas distribute data entities into multiple tables to minimize data storage redundancy and optimize data updates. Before mapping a normalized schema to the business model, you need to understand how the distributed structure can be understood in terms of facts and dimensions.

After analyzing the structure, you pick a table that has fact columns and then map the physical fact columns to one or more logical fact tables. Then, for each dimension associated with that set of physical fact columns, you map the distributed physical tables containing dimensional columns to a single logical table. Like with snowflake schemas, you can achieve this by having multiple logical table sources, or by using a single logical table source with joins. Mapping normalized schemas is an iterative process because you first map a certain set of facts, then the associated dimensions, and then you move on to the next set of facts.

Note that when a single physical table has both fact and dimension columns, you may need to create a physical alias table to handle the multiple roles played by that table.

- **Fully Denormalized Schemas.** This type of dimensional schema combines the facts and dimensions as columns in one table (or flat file), and is mapped differently than other types of schemas. When you map a fully denormalized schema to the star-shaped business model, you map the physical fact columns from the single physical fact table to multiple logical fact tables in the business model. Then, you map the physical dimension columns to the appropriate conformed logical dimension tables.

About Cubes in Multidimensional Data Sources

Cubes are made up of measures and organized by dimensions. Because they are already dimensional, each cube maps easily to the logical fact and dimension tables in the business model.

Note the following about measures and dimensions:

- Measures in multidimensional cubes and relational fact columns both map to logical measures in the Business Model and Mapping layer. However, measures in multidimensional cubes already include calculations and aggregations, unlike relational fact columns, which require the calculations and aggregations to be applied in the business model. Rather than treating cubes like relational sources, the Oracle BI Server can take advantage of the pre-aggregated data and powerful calculations in the cube.
- Multidimensional physical objects and relational physical objects both map to logical dimensions in the Business Model and Mapping layer. However, dimensional and hierarchical semantics are already built into multidimensional data sources, unlike relational sources. The Oracle BI Server can take advantage of the more complete hierarchy and dimensional support in the cube, both during import and at query time.

Identifying the Data Source Table Structure

The Administration Tool provides an interface to map logical tables to the underlying physical tables in your data sources. Before you can map the tables, you need to identify the contents of the physical data sources as it relates to your business model. To do this correctly, you need to identify the following contents of the physical data source:

- Identify the contents of each table
- Identify the detail level for each table
- Identify the table definition for each aggregate table. This lets you set up aggregate navigation. The following detail is required by the Oracle BI Server:
 - The columns by which the table is grouped (the aggregation level)
 - The type of aggregation (SUM, AVG, MIN, MAX, or COUNT)

For information on how to set up aggregate navigation in a repository, see [Chapter 10](#).

- Identify the contents of each column
- Identify how each measure is calculated
- Identify the joins defined in the database

To acquire this information about the data, you could refer to any available documentation that describes the data elements created when the data source was implemented, or you might need to spend some time with the DBA for each data source to get this information. To fully understand all the data elements, you might also need to ask people in the organization who are users of the data, owners of the data, or the application developers of applications that create the data.

Guidelines for Designing a Repository

After analyzing your business model needs and identifying the database content your business needs, you can complete your repository design. This section contains some design best practices that can help you implement a more efficient repository.

Typically, you should not make performance tuning changes until you have imported and tested your databases. These tasks are performed during the final steps in

completing the setup of your repository. For more information about these final steps, see [Chapter 14](#).

This section contains the following topics:

- [General Tips for Working on the Repository](#)
- [Design Tips for the Physical Layer](#)
- [Design Tips for the Business Model and Mapping Layer](#)
- [Design Tips for the Presentation Layer](#)

General Tips for Working on the Repository

Follow these recommended design strategies for structuring your Oracle BI repository:

- If you work in online mode, save backups of the online repository before and after every completed unit of work. If needed, use **Copy As** on the **File** menu to make an offline copy containing the changes.
- Use the Physical Diagrams in the Administration Tool to verify sources and joins.
- Decide whether you want to set up row-level security controls in the database, or in the repository. This decision determines whether you share connection pools and cache, and may limit the number of separate source databases you want to include in your deployment. See [Chapter 13, "Applying Data Access Security to Repository Objects"](#) for more information.

Most dialogs in the Administration Tool have Help that provides information about how to complete a task. To access a help topic, click the **Help** button in a dialog, or select **Help** from some menus.

Design Tips for the Physical Layer

The Physical layer contains information about the physical data sources. The most common way to create the schema in the Physical layer is by importing metadata from databases and other data sources. If you import metadata, many of the properties are configured automatically based on the information gathered during the import process. You can also define other attributes of the physical data source, such as join relationships, that might not exist in the data source metadata.

The Physical layer can contain data sources of many different types, including multidimensional, relational, and XML sources. See ["System Requirements and Certification"](#) for information about supported databases.

For each data source, there is at least one corresponding connection pool. The connection pool contains data source name (DSN) information used to connect to a data source, the number of connections allowed, timeout information, and other connectivity-related administrative details. See ["About Connection Pools"](#) for more information.

The following is a list of tips to use when designing the Physical layer:

- It is recommended that you use table aliases frequently in the Physical layer to eliminate extraneous joins, including the following:
 - Eliminate all physical joins that cross dimensions (inter-dimensional circular joins) by using aliases.
 - Eliminate all circular joins (intra-dimensional circular joins) in a logical table source in the Physical Model by creating physical table aliases.

For example, say you have a Customer table that can be used to look up ship-to addresses, and using a different join, to look up bill-to addresses. Avoid the circular joins by aliasing the table in the Physical layer so that there is one instance for each purpose, with separate joins.

If you do not eliminate circular joins, you might get erroneous report results. In addition, query performance might be negatively impacted.

- You might import some tables into the Physical layer that you might not use right away, but that you do not want to delete. To identify tables that you do want to use right away in the Business Model and Mapping layer, you can assign aliases to physical tables before mapping them to the business model layer.

Note: To have the name of a table to which you assigned an alias appear, select **Display original name for alias in diagrams** in **Tools > Options > General**.

- An opaque view (a Physical layer table that consists of a `SELECT` statement) should be used only if there is no other solution to your modeling problem. Ideally, a physical table should be created, or alternatively a materialized view. Opaque views prevent the Oracle BI Server from generating its own optimized SQL, because they contain fixed SQL statements that are sent to the underlying data source.

Design Tips for the Business Model and Mapping Layer

The Business Model and Mapping layer organizes information by business model. In this layer, each business model is effectively a separate application.

The logical schema defined in each business model must contain at least two logical tables. Relationships need to be defined between all the logical tables. See "[About Layers in the Oracle BI Repository](#)" for more information about business model schemas. See [Chapter 8](#) for more information about setting up the Business Model and Mapping layer.

The following is a list of tips to use when designing the Business Model and Mapping layer:

- Create the business model with one-to-many logical joins between logical dimension tables and the fact tables wherever possible. The business model should ideally resemble a simple star schema in which each fact table is joined directly to its dimensions.
- Every logical fact table must join to at least one logical dimension table. Note that when the source is a fully denormalized table or flat file, you must map its physical fact columns to one or more logical fact tables, and its physical dimension columns to logical dimension tables.
- Every logical dimension table should have a dimensional hierarchy associated with it. This rule holds true even if the hierarchy has only one level, such as a scenario dimension {actual, forecast, plan}.
- When creating level-based measures, make sure that all appropriate fact sources map to the appropriate level in the hierarchy using aggregation content. You set up aggregation content in the Levels tab of the Logical Column dialog for the measure. Note that this is different from the Content tab of the Logical Table Source dialog, which is used to specify the grain of the source tables to which it maps.

You only need to set up aggregation content in the Levels tab of the Logical Column dialog for level-based measures. For measures that are not level based, leave the Logical Level field blank.

- Typically, logical fact tables should not contain any keys. The only exception is when you need to send Logical SQL queries against the Oracle BI Server from a client that requires keys. In this case, you need to expose those keys in both the logical fact tables, and in the Presentation layer.
- Normally, all columns in logical fact tables are aggregated measures, except for keys required by external clients, or dummy columns used as a divider. Other non-aggregated columns should instead exist in a logical dimension table.
- In some situations, you might want to have multiple logical fact tables in a single business model. For Logical SQL queries, the multiple logical fact tables behave as if they are one table.

Reasons to have multiple logical fact tables include:

- To assign projects. See ["Setting Up Projects"](#) for more information.
- To automatically create small subject areas in the Presentation layer. See ["Automatically Creating Subject Areas Based on Logical Stars and Snowflakes"](#) for more information.
- For organization and simplicity of understanding.

Unlike relational fact tables, logical fact tables can contain measures of different grains. Because of this, grain is not a reason to split up logical fact tables.

- You can define calculations in either of the following ways:
 - Before the aggregation, in the logical table source. For example:

```
sum(col_A *( col_B))
```

- After the aggregation, in a logical column derived from two other logical columns. For example:

```
sum(col A) * sum(col B)
```

You can also define post-aggregation calculations in Answers or in Logical SQL queries.

- If you plan to use Oracle Scorecard and Strategy Management, it is a best practice to implement at least one time dimension in the Oracle BI repository you are using for your KPIs. This action is necessary because you use KPIs in scorecards to measure progress and performance over time. Note that an individual scorecard automatically picks up any dimension used by KPIs in that scorecard.
- Aggregate sources should be created as separate logical table sources. For fact aggregates, use the Content tab of the Logical Table Source dialog to assign the correct logical level to each dimension.
- Each dimension level in a hierarchy must have a unique level key. Also, each logical dimension table must have a unique primary key. Normally, this key is also used as the level key for the lowest hierarchy level.
- Renaming columns in the Business Model and Mapping layer automatically creates aliases (synonyms) for Presentation layer columns that have the property **Use Logical Column Name** selected.
- To prevent problems with aggregate navigation, ensure that each logical level of a dimension hierarchy contains the correct value in the field named **Number of**

elements at this level. Fact sources are selected on a combination of the fields selected as well as the levels in the dimensions to which they map. By adjusting these values, you can alter the fact source selected by the Oracle BI Server. See ["Creating Logical Levels in a Dimension"](#) for more information about setting this value.

Modeling Outer Joins

The following guidelines provide tips on how to model outer joins:

- Due to the nature of outer joins, queries that use them are usually slower. Because of this, define outer joins only when necessary. Where possible, use ETL techniques to eliminate the need for outer joins in the reporting SQL.
- Outer joins are always defined in the Business Model and Mapping layer. Physical layer joins do not specify inner or outer.
- You can define outer joins by using logical table joins, or in logical table sources. Which type of outer join you use is determined by whether the physical join maps to a business model join, or to a logical table source join.
- Be aware that outer joins in logical table sources are always included in a query, even if the none of the columns in one of the mapped physical tables are used.

For example, assume that a logical table source is mapped to physical tables A and B. When no outer joins are defined, if physical table A is not required to satisfy a request, it is not included in the physical query. However, if the logical table source has an outer join defined to table A, table A is still included in the physical query, even if only physical table B is required to satisfy the request.

- If you must define an outer join, try to create two separate dimensions, one that uses the outer join and one that does not. Make sure to name the dimension with the outer join in a way that clearly identifies it, so that client users can use it as little as possible.

Design Tips for the Presentation Layer

You set up the user view of a business model in the Presentation layer. The names of folders and columns in the Presentation layer can appear in localized language translations. The Presentation layer is the appropriate layer in which to set user permissions. For complete information about working in the Presentation layer, see [Chapter 11](#).

In this layer, you can do the following:

- You can show fewer columns than exist in the Business Model and Mapping layer. For example, you can exclude the key columns because they have no business meaning.
- You can organize columns using a different structure from the table structure in the Business Model and Mapping layer.
- You can display column names that are different from the column names in the Business Model and Mapping layer.
- You can set permissions to grant or deny users access to individual subject areas, tables, and columns.
- You can export logical keys to ODBC-based query and reporting tools.
- You can create multiple subject areas for a single business model.

- You can create a list of aliases (synonyms) for presentation objects that can be used in Logical SQL queries. This feature lets you change presentation column names without breaking existing reports.

The following is a list of tips to use when designing the Presentation layer:

- Because there is no automatic way to synchronize all changes between the Business Model and Mapping layer and the Presentation layer, it is best to wait until the Business Model and Mapping layer is relatively stable before adding customizations in the Presentation layer.
- There are many ways to create subject areas, such as dragging and dropping the entire business model, dragging and dropping incremental pieces of the model, or automatically creating subject areas based on logical stars or snowflakes. See ["Creating Subject Areas"](#) for information about each of these methods. Dragging and dropping incrementally works well if certain parts of your business model are still changing.
- It is a best practice to rename objects in the Business Model and Mapping layer rather than the Presentation layer, for better maintainability. Giving user-friendly names to logical objects rather than presentation objects ensures that the names can be reused in multiple subject areas. Also, it ensures that the names persist even when you need to delete and re-create subject areas to incorporate changes to your business model.
- Be aware that members in a presentation hierarchy are not visible in the Presentation layer. Instead, you can see hierarchy members in Answers.
- You can use the Administration Tool to update Presentation layer metadata to give the appearance of nested folders in Answers. See ["Nesting Folders in Answers"](#) for more information.
- When setting up data access security for a large number of objects, consider setting object permissions by role rather than setting permissions for individual columns. See [Chapter 13, "Applying Data Access Security to Repository Objects"](#) for details.
- When setting permissions on presentation objects, you can change the default permission by setting the `DEFAULT_PRIVILEGES` configuration setting in the `NQConfig.INI` file. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Topics of Interest in Other Guides

Some topics that may be of interest to metadata repository builders are covered in other guides. [Table 1–1](#) lists these topics, and indicates where to go for more information.

Table 1–1 Topics Covered in Other Guides

Topic	Where to Go for More Information
Starting and stopping Oracle Business Intelligence processes	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Using the Oracle BI Server XML API to work with your repository	<i>Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Using the Oracle BI Server Web services	<i>Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition</i>

Table 1–1 (Cont.) Topics Covered in Other Guides

Topic	Where to Go for More Information
Setting up and managing query caching	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing configuration settings that affect repository development in Fusion Middleware Control and NQSCONFIG.INI	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing users, groups, and application roles	<i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Moving from test to production environments	<i>Oracle Fusion Middleware Administrator's Guide</i>
Setting up DSNs for the Oracle BI Server	<i>Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Localizing Oracle Business Intelligence deployments	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Information about the SA System subject area	<i>Oracle Fusion Middleware Scheduling Jobs Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing logging	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing usage tracking	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
General information about managing Oracle WebLogic Server	<i>Oracle Fusion Middleware Administrator's Guide</i>

System Requirements and Certification

Refer to the system requirements and certification documentation for information about hardware and software requirements, platforms, databases, and other information. Both of these documents are available on Oracle Technology Network (OTN).

The system requirements document covers information such as hardware and software requirements, minimum disk space and memory requirements, and required system libraries, packages, or patches:

http://www.oracle.com/technology/software/products/ias/files/fusion_requirements.htm

The certification document covers supported installation types, platforms, operating systems, databases, JDKs, and third-party products:

http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html

Before You Begin

This chapter provides an overview of the Administration Tool, and explains other concepts that you need to know before beginning to build your metadata repository.

This chapter contains the following topics:

- [About the Oracle BI Administration Tool](#)
- [About the Oracle BI Server Command-Line Utilities](#)
- [About Options in Fusion Middleware Control and NQSCONFIG.INI](#)
- [About the SampleApp.rpd Demonstration Repository](#)
- [Using Online and Offline Repository Modes](#)
- [Checking the Consistency of a Repository or a Business Model](#)

About the Oracle BI Administration Tool

The Oracle BI Administration Tool is a Windows application that you can use to create and edit repositories.

This section describes the Administration Tool main window, how to set preferences, Administration Tool menus, and other related information.

This section contains the following topics:

- [Opening the Administration Tool](#)
- [About the Administration Tool Main Window](#)
- [Setting Preferences](#)
- [About Administration Tool Menus](#)
- [Using the Physical and Business Model Diagrams](#)
- [Editing, Deleting, and Reordering Objects in the Repository](#)
- [About Naming Requirements for Repository Objects](#)
- [Using the Browse Dialog to Browse for Objects](#)
- [Changing Icons for Repository Objects](#)
- [Sorting Objects in the Administration Tool](#)
- [About Features and Options for Oracle Marketing Segmentation](#)

Opening the Administration Tool

To open the Administration Tool, choose **Start > Programs > Oracle Business Intelligence > BI Administration**.

Note: Do not open the Administration Tool by double-clicking a repository file. The resulting Administration Tool window is not initialized to your Oracle instance, and errors will result.

You can also launch the Administration Tool from the command line, as follows:

1. In Windows Explorer, go to the following location:
`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup`
2. Double-click `bi-init.cmd`.
3. Type `admintool` and press Enter.

About the Administration Tool Main Window

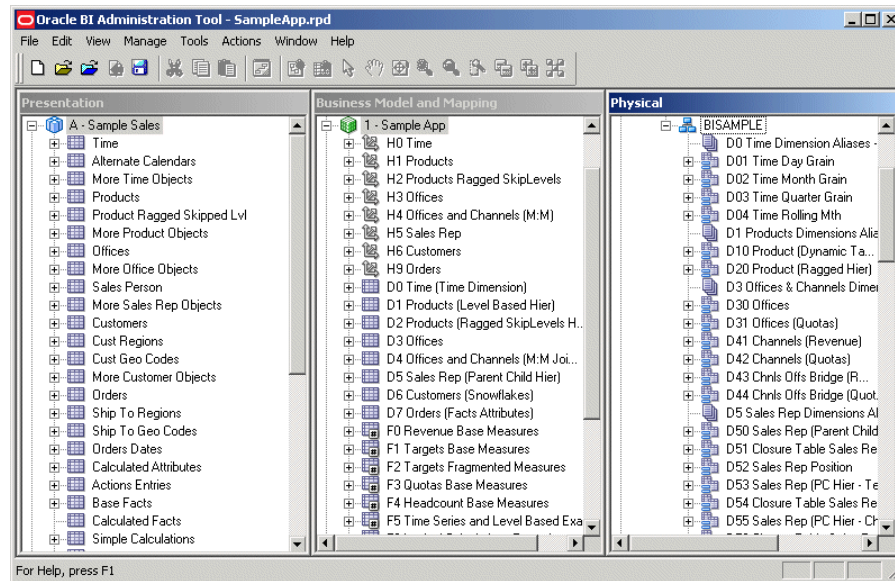
The main window of the Administration Tool shows a graphical representation of the three layers of a repository (the Physical layer, Business Model and Mapping layer, and Presentation layer). See "[About Layers in the Oracle BI Repository](#)" for more information.

The Administration Tool main window also contains the following:

- **Menus.** See "[About Administration Tool Menus](#)" for more information.
- **Toolbar.** Provides access to global functionality such as Open and Save, and also includes functions for the Physical Diagram and Business Model Diagram.
- **Status bar.** Provides contextual information about the current dialog or selected object, as well as other useful information.
- **Title bar.** In offline mode, displays the name of the open repository. In online mode, displays the DSN for the Oracle BI Server to which you are connected.

[Figure 2–1](#) shows the Administration Tool main window.

Figure 2–1 Example Administration Tool Main Window



Setting Preferences

You can use the Options dialog to set preferences for the Administration Tool.

To set preferences:

1. In the Administration Tool, select **Tools**, then select **Options** to display the Options dialog.
2. On the General tab, select the options you want to choose.

Table 2–1 describes the options on the General tab.

Table 2–1 Options on the General Tab

Option	Action When Selected
Tile when resizing	Automatically tiles the layer panes of the repository when you resize the Administration Tool. When this option is selected, the Cascade and Tile options are not available in the Windows menu of the Administration Tool.
Display qualified names in diagrams	<p>Displays fully qualified names in the Physical Diagram and Business Model Diagram. For example, selecting this option displays "B - Sample Fcst Data"..."B02 Market" rather than B02 Market in the Physical Diagram.</p> <p>Selecting this option can help identify objects by including the name of the parent database or business model, but it can also make the diagram harder to read because the fully qualified names are longer.</p> <p>Note: If you choose not to select this option, you can still see fully qualified names by moving the cursor over an object in the diagram, or by selecting an object in the diagram and then viewing the text in the status bar.</p>
Display original names for alias in diagrams	Displays the names of original physical tables rather than the names of alias tables in the Physical diagram. Select this option when you want to identify the original table rather than the alias table name.

Table 2–1 (Cont.) Options on the General Tab

Option	Action When Selected
Show Calculation Wizard introduction page	<p>Displays the Calculation Wizard introduction page. The introduction page also contains an option to suppress its display in the future.</p> <p>Use the Calculation Wizard to create new calculation columns that compare two existing columns and to create metrics in bulk (aggregated), including existing error trapping for NULL and divide by zero logic. See "Using the Calculation Wizard" for more information.</p>
Check out objects automatically	<p>Automatically checks out an object when you double-click it. If you do not select this option, you are prompted to check out objects before you can edit them.</p> <p>This option only applies when the Administration Tool is open in online mode. See "Opening a Repository in Online Mode" for more information.</p>
Show row count in physical view	<p>Displays row counts for physical tables and columns in the Physical layer. Row counts are not initially displayed until they are updated. To update the counts, select Tools > Update All Row Counts. You can also right-click a table or column in the Physical layer and select the option Update Row Count.</p> <p>Note: Row counts are not shown for items that are stored procedure calls (from the Table Type list in the General tab of the Physical Table dialog). Row counts are not available for XML, XML Server, or multidimensional data sources. When you are working in online mode, you cannot update row counts on any new objects until you check them in.</p>
Show toolbar	When selected, displays the Administration Tool toolbar.
Show statusbar	When selected, displays the Administration Tool status bar.
Prompt when moving logical columns	Lets you ignore, specify an existing, or create a new logical table source for a moved column.
Remove unused physical tables after Merge	Executes a utility to clean the repository of unused physical objects. It might make the resulting repository smaller.
Allow import from repository	<p>When selected, the Import from Repository option on the File menu becomes available.</p> <p>Note: By default, the Import from Repository option on the File menu is disabled and this option will not be supported in the future. It is recommended that you create projects in the repository that contain the objects that you want to import, and then use repository merge to bring the projects into your current repository. See "Merging Repositories" for more information.</p>
Allow logical foreign key join creation	When selected, provides the capability to create logical foreign key joins with the Joins Manager. This option is provided for compatibility with previous releases and is generally not recommended.
Skip Gen 1 levels in Essbase drag and drop actions	<p>When selected, excludes Gen 1 levels when you drag and drop Essbase cubes or dimensions from the Physical layer to the Business Model and Mapping layer. Often, Gen 1 levels are not needed for analysis, so they can be excluded from the business model.</p> <p>See "Working with Essbase Data Sources" for more information.</p>

Table 2–1 (Cont.) Options on the General Tab

Option	Action When Selected
Hide unusable logical table sources in Replace wizard	<p>By default, the Replace Wizard shows all logical table sources, even ones that are not valid for replacement. When this option is selected, unusable logical table sources are hidden in the Replace Wizard screens. Click Info for details on why a logical table source that maps to that column does not appear in the list.</p> <p>Selecting this option might result in the Wizard page loading more quickly, especially for large repositories.</p>
Allow first Connection Pool for Init Blocks	<p>By default, when you select a connection pool for an initialization block, the first connection pool under the database object in the Physical layer does not show up as available for selection. This behavior ensures that you cannot use the same connection pool for initialization blocks that you use for queries. If the same connection pool is used for initialization blocks and for queries, then queries might be blocked whenever initialization blocks run. Alternatively, initialization blocks used for authentication might be blocked by long-running queries, causing delayed or hanging logins.</p> <p>Select this option to change the default behavior and allow the first connection pool to be selected for initialization blocks. Note that selecting this option is not a best practice and might cause performance issues.</p> <p>See "About Connection Pools for Initialization Blocks" for more information.</p>

3. On the Repository tab, you can set the following options:

- **Show tables and dimensions only under display folders.** You can create display folders to organize objects in the Physical and Business Model and Mapping layers. They have no metadata meaning. After you create a display folder, the selected objects appear in the folder as a shortcut and in the database or business model tree as an object. You can hide the objects so that only the shortcuts appear in the display folder.

See "[Setting Up Display Folders in the Physical Layer](#)" and "[Setting Up Display Folders in the Business Model and Mapping Layer](#)" for more information about creating display folders.
- **Hide level based measure.** By default, each level of a dimension hierarchy in the Business Model and Mapping layer shows both dimension columns that are assigned to that level, and level-based measures that have been fixed at that level. Level-based measures are objects that are not part of the dimension table, but that have been explicitly defined as being at a particular level.

Hiding level-based measures in dimension hierarchies can reduce clutter. Note that the measures are still visible in the logical fact tables.

See [Example 9–1, "Level-Based Measure Calculations"](#) for more information about level-based measures.
- **Default logging level.** This option determines the default query logging level for the internal BISystem user. The BISystem user owns the Oracle BI Server system processes and is not exposed in any user interface.

A query logging level of 0 (the default) means no logging. Set this logging level to 2 to enable query logging for internal system processes like event polling and initialization blocks.

See "Managing the Query Log" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about the query log and query logging levels.

- **LDAP.** If you are using any alternative LDAP servers, the Oracle BI Server maintains an authentication cache in memory for user identifiers and properties, which improves performance when using LDAP to authenticate large numbers of users. Disabling the authentication cache can slow performance when hundreds of sessions are being authenticated. Note that the authentication cache is not used for Oracle WebLogic Server's embedded directory server.

Properties for the authentication cache include:

- **Cache refresh interval.** The interval at which the authentication cache entry for a logged on user is refreshed.
- **Number of Cache Entries.** The maximum number of entries in the authentication cache, preallocated when the Oracle BI Server starts. If the number of users exceeds this limit, cache entries are replaced using the LRU algorithm. If this value is 0, then the authentication cache is disabled.

You need to specify some additional LDAP properties when you are using a secure connection to your LDAP server. In other words, provide the following information when you have selected **SSL** on the Advanced tab of the LDAP Server dialog:

- **Key file name.** The name of the key file that holds the client certificate and Certificate Authority (CA) certificate.
- **Password and Confirm password.** The password for the key file.

Note that the authentication cache properties and key file properties are shared for all defined LDAP server objects.

4. On the Sort Objects tab, specify which repository objects appear in the Administration Tool in alphabetical order. For example, if you want the database objects that appear in the Physical layer to appear in alphabetical order, select the **Database** option.
5. On the Cache Manager tab, select the columns you want to display in the Cache Manager. To change the order of columns in the Cache Manager, select an item, then use the **Up** and **Down** buttons to change its position.
6. On the Multiuser tab, specify the path to the multiuser development directory and the name of the local developer for this Administration Tool. See "[Setting Up a Pointer to the Multiuser Development Directory](#)" for more information.
7. On the More tab, you can set the scrolling speed for Administration Tool dialogs. To set the scrolling speed, position the cursor on the slider.
8. Click **OK** when you are finished setting preferences.

About Administration Tool Menus

The Administration Tool includes menus for File, Edit, View, Manage, Tools, Actions, Window, and Help. These menus are described in the following sections.

File Menu

The File menu provides options to work with repositories, like **Open** and **Save**, as well as several server-related options like **Check Out All** that are only active when a

repository is open in online mode. The File menu also provides a list of recently opened files.

Table 2–2 lists the options in the File menu.

Table 2–2 File Menu Options

Menu Option	Description
New Repository	Opens the Create New Repository Wizard and closes the currently open repository, if any. If a repository is currently open with unsaved changes, you are prompted to save them before proceeding. See "Creating New Oracle BI Repository Files" for more information.
Open	Provides options for opening a repository in either offline or online mode. See "Using Online and Offline Repository Modes" for more information.
Multiusers	Provides options to check out projects in a multiuser development environment and view multiuser development history. See Chapter 3, "Setting Up and Using the Multiuser Development Environment" for more information.
Close	Closes the currently open repository. If you have unsaved changes, you are prompted to save them.
Save	Saves your latest changes.
Save As	Opens the Save As dialog so that you can save the repository to a different file. The new file remains open in the Administration Tool.
Copy As	Opens the Save Copy As dialog so that you can copy the repository to a different file. The current file, not the new file, remains open in the Administration Tool.
Change Password	Lets you change the repository password for the currently open repository. See "Changing the Repository Password" for more information.
Import Metadata	Opens the Import Metadata Wizard. See the following sections for more information: <ul style="list-style-type: none"> ■ "Importing Metadata from Relational Data Sources" ■ "Importing Metadata from Multidimensional Data Sources" ■ "Working with ADF Business Component Data Sources" ■ "Importing Metadata from XML Data Sources"
Compare	Prompts you to select the repository with which you want to compare the currently open repository and opens the Compare repositories dialog. See "Comparing Repositories" for more information.
Turn off Compare Mode	Turns off any highlighted changed objects. This option is only available if you have turned on compare mode by choosing Mark in the Compare repositories dialog.
Merge	Opens the Merge Repository Wizard. See "Merging Repositories" for more information.

Table 2–2 (Cont.) File Menu Options

Menu Option	Description
Check Global Consistency	Checks the repository for consistency and opens the Consistency Check Manager. See "Checking the Consistency of a Repository or a Business Model" for more information.
Check Out All	Checks out all repository objects. This option is only available in online mode.
Check In Changes	Checks in all repository objects. This option is only available in online mode.
Undo All Changes	Rolls back all changes made since the last check-in. This option is only available in online mode.
Exit	Closes the currently open repository and then closes the Administration Tool. If you have unsaved changes, you are prompted to save them.

Edit Menu

The Edit menu provides access to the following basic editing functions for repository objects: **Cut**, **Copy**, **Paste**, **Duplicate**, and **Delete**. You can also choose **Properties** to view and edit properties for a selected object.

View Menu

The View menu options let you hide or display the panes that show the three layers of the repository (Presentation, Business Model and Mapping, and Physical). You can also display the Business Model Diagram and Physical Diagram.

Choose **Refresh** to refresh the repository view. This feature can be useful in online mode to reveal changes made by other clients. It can also be used in either online or offline mode when the repository view has become out of sync and does not display a recent change or addition. Refreshing the repository view collapses any expanded objects in the tree panes and helps reduce clutter.

Manage Menu

The Manage menu enables you to access the management functions described in [Table 2–3](#).

Table 2–3 Manage Menu Options

Menu Option	Description
Jobs	Opens the Job Manager. The Job Manager is the management interface to Oracle BI Scheduler. This option is available when a repository is open in online mode. See "Using Oracle BI Scheduler Job Manager" in <i>Oracle Fusion Middleware Scheduling Jobs Guide for Oracle Business Intelligence Enterprise Edition</i> for more information.
Sessions	Opens the Session Manager. In the Session Manager, you can monitor activity on the system, including the current values of repository and session variables. This option is available when a repository is open in online mode. See "Managing Server Sessions" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i> for more information.

Table 2–3 (Cont.) Manage Menu Options

Menu Option	Description
Cache	<p>Opens the Cache Manager. The Cache Manager enables you to monitor and manage the cache. This option is available when a repository is open in online mode and caching is enabled.</p> <p>See "Using the Cache Manager" in <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i> for more information about enabling the cache and using the Cache Manager.</p>
Clusters	<p>Opens the Cluster Manager. The Cluster Manager monitors and manages the operations and activities of the cluster. This option is available when the Oracle BI Cluster Server is installed.</p> <p>See "Using the Cluster Manager" in <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i> for more information.</p>
Identity	<p>Opens the Identity Manager. The Identity Manager provides access to data access security functions and other identity-related options.</p> <p>See Chapter 13, "Applying Data Access Security to Repository Objects" for more information.</p>
Joins	<p>Opens the Joins Manager. The Joins Manager enables you to work with both physical and logical joins in a list format. The Joins Manager provides an alternative to working with joins in the Physical and Business Model Diagrams and shows all join types in one place.</p> <p>See "Defining Physical Joins with the Joins Manager" and "Defining Logical Joins with the Joins Manager" for more information.</p>
Variables	<p>Opens the Variable Manager. The Variable Manager enables you to create, edit, or delete variables and initialization blocks.</p> <p>See Chapter 18, "Using Variables in the Oracle BI Repository" for more information.</p>
Projects	<p>Opens the Project Manager. The Project Manager enables you to create, edit, or remove projects or project elements. Project elements include subject areas (formerly called presentation catalogs), logical fact tables, groups, users, variables, and initialization blocks. You use projects during multiuser development.</p> <p>See "Setting Up Projects" for more information.</p>
Marketing	<p>Applies to the Oracle Marketing Segmentation product. For information about using the Marketing options in Oracle Business Intelligence, see <i>Oracle Marketing Segmentation Guide</i>.</p>

Tools Menu

The Tools menu options enable you to access the functions described in [Table 2–4](#).

Table 2–4 Tools Menu Options

Menu Option	Description
Update All Row Counts	<p>Updates row counts in the Physical layer.</p> <p>See "Displaying and Updating Row Counts for Physical Tables and Columns" for more information.</p>
Show Consistency Checker	<p>Opens the Consistency Check Manager.</p> <p>See "Checking the Consistency of a Repository or a Business Model" for more information.</p>

Table 2–4 (Cont.) Tools Menu Options

Menu Option	Description
Query Repository	Opens the Query Repository dialog. See "Querying and Managing Repository Metadata" for more information.
Utilities	Opens the Utilities dialog, which lets you select from a list of Administration Tool utilities. See "Using Administration Tool Utilities" for more information.
Options	Opens the Options dialog, which lets you customize Administration Tool display preferences and other options. See "Setting Preferences" for more information.

Actions Menu

The Actions menu options are available when working with the Physical Diagram or Business Model Diagram. The options enable you to select elements, create new joins, create new tables, and perform other diagram operations. Every toolbar option for the diagrams has an Actions menu equivalent.

Window Menu

The Window menu options enable you to cascade or tile open layer windows and toggle among them.

Help Menu

The Help menu provides the following options:

- **Help Topics.** Access the Help system for the Administration Tool.
- **Oracle BI on the Web.** Access the Oracle Business Intelligence home page on the Oracle Technology Network (OTN).
- **About Oracle BI Administration Tool.** Obtain version information about the Administration Tool.

Using the Physical and Business Model Diagrams

You can use the Physical and Business Model Diagrams in the Administration Tool to see a graphical view of physical and logical tables and joins. You can choose to view tables in expanded mode, with columns visible, or in collapsed mode, where only the name of the table is displayed. This section describes the layout and navigation capabilities for both diagrams.

After launching the Physical or Business Model Diagram, you can use toolbar options to zoom, pan, and control the layout of the tables. [Table 2–5](#) describes the available toolbar options.

Table 2–5 Toolbar Options for the Physical and Business Model Diagrams


Option Name	Description
Auto Layout 	Select this option to revert to an automatically assigned symmetric table layout. Any customizations you have made to the layout (by manually moving individual tables) will be lost.

Table 2–5 (Cont.) Toolbar Options for the Physical and Business Model Diagrams













Option Name	Description
<p>Expand All</p> 	<p>Select this option to show all tables in expanded view, with columns showing. Tables in expanded view appear like the following:</p>  <p>Note the following additional features for expanded tables:</p> <ul style="list-style-type: none"> ■ Use the scrollbar to scroll down the full list of columns. ■ Click a column heading to sort based on that column. ■ Double-click a table in expanded view to launch the Properties dialog for that object. ■ Click the Collapse icon in the upper right corner to collapse an individual table object. ■ To resize expanded tables, select a table, mouse over a handle, and then click and drag the handle.
<p>Collapse All</p> 	<p>Select this option to show all tables in collapsed view, with only the table name showing. Tables in collapsed view appear like the following:</p>  <p>You can double-click an individual table in collapsed view to expand only that object.</p>
<p>Marquee Zoom</p> 	<p>Select this option to use the Marquee Zoom tool, which lets you select a particular region to which you want to zoom. To use Marquee Zoom, left-click, hold, and drag to define a rectangular region where you want to zoom.</p>
<p>Zoom Out</p> 	<p>Select this option to cause the diagram view to zoom out one level.</p>
<p>Zoom In</p> 	<p>Select this option to cause the diagram view to zoom in one level.</p>
<p>Fit</p> 	<p>Select this option to cause the layout to dynamically adjust to the current diagram window size so that all objects fit in the window.</p>
<p>Pan</p> 	<p>Select this option to use the Pan tool, which lets you pan around the current layout. Left-click, hold, and drag to move the view.</p> <p>This option is especially useful when the diagram layout exceeds the available space.</p>

Table 2–5 (Cont.) Toolbar Options for the Physical and Business Model Diagrams

Option Name	Description
Select 	<p>Select this option to enable the ability to select objects in the diagram. You can double-click a join or expanded table object to access the Properties dialog, or you can select a particular table and drag it to a new location. Note that location information is not saved after you close the diagram or choose Auto Layout.</p> <p>You can select multiple objects using the SHIFT or CTRL keys. Press SHIFT and select multiple objects, or click and drag to define an area where you want all objects selected. Press CTRL to individually add or remove particular objects to the selection set.</p>
New Table 	<p>Select this option to create a new physical or logical table while in the diagram view. Left-click the background to launch the Properties dialog for the new object, and then provide details as necessary. For physical tables, you first need to select the parent object under which the new table will be created (such as a schema, catalog, or database object).</p> <p>See also the following sections:</p> <ul style="list-style-type: none"> ■ "Creating and Managing Physical Tables and Physical Cube Tables" ■ "Creating Logical Tables"
New Join 	<p>Select this option to create a new join while in the diagram view. First, left-click the first table in the join (the table representing many in the one-to-many join). Then, move the cursor to the table to which you want to join (the table representing one in the one-to-many join), and then left-click the second table to select it. Provide details in the Properties dialog for the new object as necessary.</p> <p>Joins in the Physical and Business Model Diagrams are represented by a line with an arrow at the "one" end of the join. Note that this display is different from the line with crow's feet at the "many" end of the join that was used in previous releases.</p> <p>See also the following sections:</p> <ul style="list-style-type: none"> ■ "Defining Physical Joins with the Physical Diagram" ■ "Defining Logical Joins with the Business Model Diagram"

Note the following additional features of the Physical and Business Model Diagrams:

- All toolbar options for the diagram, such as Select, New Table, and New Join, are also available from the Actions menu.
- Moving the mouse over a table causes the fully-qualified name for that table to appear in the status bar.
- You can have both the Physical Diagram and Business Model Diagram windows open at the same time.
- Any customizations you have made to the layout (by manually moving individual tables) are lost after you close the diagram or choose Auto Layout.
- You can cause fully-qualified table names to appear in diagrams by setting a preference in the Options dialog. See ["Setting Preferences"](#) for more information.

See also the following sections for more information about using the Physical and Business Model Diagrams:

- ["Physical Diagram and Business Model Diagram Keyboard Shortcuts"](#)

- ["Working with the Physical Diagram"](#)
- ["Working with the Business Model Diagram"](#)

Editing, Deleting, and Reordering Objects in the Repository

This section provides information about editing, deleting, and reordering objects.

- To edit objects, double-click an object, or right-click an object and select **Properties**. Then, complete the fields in the dialog that is displayed. In some dialogs, you can click **Edit** to open the appropriate dialog.
- To delete objects, select one or more objects and click **Delete**, or press the delete key. You can also right-click an object and select **Delete**.
- To reorder objects, drag and drop an object to a new location. Note the following:
 - Reordering is only possible for certain objects and in certain dialogs.
 - In some dialogs, you can use an up or down arrow to move objects to a new location.
 - In the Administration Tool main window, you can drag and drop an object onto its parent to duplicate the object. For top-level objects like business models and subject areas, drag and drop the object onto white space to duplicate it.

About Naming Requirements for Repository Objects

All repository object names must follow these requirements:

- Names cannot be longer than 128 characters
- Names cannot contain leading or trailing spaces
- Names cannot contain single quotes, question marks, or asterisks

Note that repository object names can include multibyte characters.

Using the Browse Dialog to Browse for Objects

The Browse dialog appears in many situations in the Administration Tool. You use it to find and select an object.

The Browse dialog is accessible from several dialogs that let you make a selection from among existing objects.

The left pane of the Browse dialog lets you browse the tree view for a particular object. It contains the following parts:

- A tree listing all of the objects in the Presentation layer, Business Model and Mapping layer, or the Physical layer of a repository.
- Tabs at the bottom of the left pane let you select a layer. Some tabs might not appear if objects from those layers are not appropriate for the task you are performing.

The right pane of the Browse dialog lets you search for the object you want. It contains the following parts:

- **Query** enables you to query objects in the repository by name and type. The **Name** field accepts an asterisk (*) as the wildcard character, so you can query for partial matches.

- The **Show Qualified Names** option lets you identify to which parents an object belongs.
- **View** lets you view properties of a selected object in read-only mode.

Note that in general, the left pane and the right pane of the Browse dialog are not connected. Rather, the panes provide alternate methods to locate the object you want.

The exception to this is the Synchronize Contents feature, which lets you synchronize an object from the query results list with the tree view. This feature is a helpful contextual tool that locates a particular object in the tree view.

Table 2–6 lists and describes the tasks you can perform in the Browse dialog.

Table 2–6 Tasks You Can Perform in the Browse Dialog

Task	Description
Querying for an object	<p>Follow these steps to query for an object:</p> <ol style="list-style-type: none"> 1. Select the object type from the Type list. 2. Type the name of the object, or a part of the name and the wildcard character (*), in the Name field. For example: <ul style="list-style-type: none"> - To search for logical tables that have names beginning with the letter Q, select Logical Tables from the Type list, and then type Q* in the Name field. - To search for logical tables that have names ending with the letters dim, type *dim in the name field. 3. Click Query. <p>Relevant objects appear in the query results list.</p>
Selecting an object	<p>Use the tree view in the left pane or the filtered view in the right pane to locate the object you want, then double-click the object.</p> <p>The Browse dialog closes, and the object is displayed in the previous dialog.</p>
Synchronizing an object in the query results list with the tree view	<p>Select an object in the Query list and then click the Synchronize Contents button.</p> <p>The object you selected is highlighted in the tree view in the left pane.</p>
Finding multiple occurrences of an object in the tree view	<p>Select an object in the tree view, such as a logical column, then click the down arrow button.</p> <p>The next occurrence of that object is highlighted in the tree view.</p>

Changing Icons for Repository Objects

In the Administration Tool, you can change the icon that represents a particular object in the repository. Changing the icon for a particular object does not have any functional effect, and is not visible in Answers or other clients. This feature is intended as a useful way to visually distinguish objects for the convenience of repository developers.

For example:

- You can use a special icon for objects that are in the Business Model and Mapping layer, but not the Presentation layer, for easier maintenance of the repository.
- You can mark objects that are logical calculations with a separate icon.
- You can choose an icon to visually distinguish tables in the Presentation layer that appear as nested folders in Answers.

- You can use an icon to denote objects in a logical table that pertain to a specific functional area, or that are sourced from a particular logical table source.

You can only change the icon for individual objects. You cannot globally change the icon for all objects of a particular type.

To change the icon for a particular repository object:

1. In the Administration Tool, right-click an object in the Physical, Business Model and Mapping, or Presentation layer (for example, a particular logical table).
2. Select **Set Icon**.
3. In the Select Icon dialog, select the icon you want to use for that object and click **OK**.

Sorting Objects in the Administration Tool

Many dialogs in the Administration Tool show lists of objects, such as a list of physical columns in the Physical Table dialog, a list of logical levels for Preferred Drill Path in the Logical Level dialog, and a list of presentation hierarchies in the Presentation Table dialog.

You can click the header to sort the objects in ascending or descending order. An up arrow or down arrow icon is displayed next to the header name, indicating how the list has been sorted.

Each list also has a default order that is persisted from session to session. The default order appears when you view a list in a dialog for the first time each session. The default order is displayed when there is no ascending or descending arrow icon in the header. Click the header three times to toggle between ascending, descending, and default order. In some cases, the default order is the ascending or descending order.

Some dialogs provide the capability to move items up or down in a list. In these dialogs, if you click **Up** or **Down** while the list is sorted in ascending or descending order, the selected item moves, and the resulting order becomes the new default order. Note that clicking the header eliminates any manually determined order.

About Features and Options for Oracle Marketing Segmentation

Some features and options in the Administration Tool are for use by organizations that have the Oracle Marketing Segmentation product. For information about these features and options, see *Oracle Marketing Segmentation Guide*.

Note that additional information about Oracle Marketing Segmentation features is provided in the Presentation Services Help.

About the Oracle BI Server Command-Line Utilities

You can use a variety of command-line utilities with the Oracle BI Server to make programmatic changes to your repository file, run sample queries, delete unwanted repository objects, and perform other tasks.

[Table 2-7](#) describes the Oracle BI Server command-line utilities.

Table 2–7 Oracle BI Server Command-Line Utilities

Utility Name	Description	Where to Go for More Information
XML utilities (biserverxmlgen, biserverxmlxec, biserverxmlcli)	Used to leverage the Oracle BI Server XML API for metadata migration, programmatic metadata generation and manipulation, metadata patching, and other functions. The XML utilities include: <ul style="list-style-type: none"> ▪ biserverxmlgen: generates XML from an existing RPD ▪ biserverxmlxec: executes the XML in offline mode to create or modify a repository file ▪ biserverxmlcli: executes the XML against the Oracle BI Server 	"About the Oracle BI Server XML API" in <i>Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
nqcmd	Used to run test queries against the repository. Connects using an Oracle BI Server ODBC DSN.	"Testing and Refining the Repository"
patchrpd	Used to generate and apply an XML patch file. This utility is especially useful for patching repository files on Linux or UNIX systems.	"Using patchrpd to Apply a Patch"
extractprojects	Used to extract projects from a given repository.	"Using the extractprojects Utility to Extract Projects"
equalizerpds	Used to equalize objects in two repositories that have the same name, but different upgrade IDs. Running this utility before merging repositories prevents unintended renaming during the merge.	"Equalizing Objects"
prunerpd	Used to delete unwanted repository objects from your repository file, such as databases, tables, columns, initialization blocks, and variables.	"Deleting Unwanted Objects from the Repository"
validaterpd	Used to check the consistency of a repository.	"Using the validaterpd Utility to Check Repository Consistency"
sametaexport	Used to generate the information necessary for the Oracle Database SQL Access Advisor or IBM DB2 Cube Views tool to preaggregate relational data and improve query performance.	"Exchanging Metadata with Databases to Enhance Query Performance"

Running bi-init to Launch a Shell Window Initialized to Your Oracle Instance

Before running any of the Oracle BI Server command-line utilities, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX) to launch a command prompt or shell window that is initialized to your Oracle instance. You can find this utility in:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup`

Then, run the appropriate command-line tool from the resulting shell window with the desired options.

About Options in Fusion Middleware Control and NQSConfig.INI

Many configuration settings that affect the Administration Tool and repository development are managed in either Fusion Middleware Control, or the NQSConfig.INI configuration file. Repository developers must be familiar with Fusion Middleware Control and NQSConfig.INI configuration settings to effectively work with the Administration Tool and with their repositories.

Some of the most common configuration settings that affect repository development include:

- **Repository File:** This option is set in Fusion Middleware Control. It controls the current published repository.
- **Disallow RPD Updates:** This option is set in Fusion Middleware Control. It controls whether the Administration Tool opens in read-only mode, in both offline and online mode.
- **LOCALE:** This option is set in NQSConfig.INI. It specifies the locale in which data is returned from the server and determines the localized names of days and months.
- **DATE_TIME_DISPLAY_FORMAT, DATE_DISPLAY_FORMAT, TIME_DISPLAY_FORMAT:** These options are set in NQSConfig.INI. They control the display of date/time formats.
- **DEFAULT_PRIVILEGES:** This option is set in NQSConfig.INI. It determines the default privilege (NONE or READ) granted to users and application roles for repository objects without explicit permissions set.

See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for full information about Fusion Middleware Control and NQSConfig.INI configuration settings.

About the SampleApp.rpd Demonstration Repository

Oracle Business Intelligence provides a sample repository called SampleApp.rpd that provides best practices for modeling many different types of objects described in this guide.

A basic version of SampleApp.rpd, called SampleAppLite.rpd, is automatically installed as the default repository when you choose the Simple Install option. SampleAppLite.rpd is located in the following directory:

```
ORACLE_INSTANCE\bifoundation\OracleBIServerComponent\coreapplication_
obisn\repository
```

The full version of SampleApp.rpd contains many additional examples and features. This version can be found on the Oracle Technology Network at:

<http://oracle.com/technetwork/middleware/bi-foundation/obiee-samples-167534.html>

The default password for SampleAppLite.rpd is Admin123. For security reasons, you must immediately change this default password the first time you open SampleAppLite.rpd in the Administration Tool. See "[Changing the Repository Password](#)" for more information.

Note: The paint repository shipped with previous versions of Oracle Business Intelligence is also available with this release in the following directory:

```
ORACLE_  
INSTANCE\bifoundation\OracleBIServerComponent\coreapplication_  
obisn\sample\paint
```

Paint.rpd is not automatically set up as the default repository. In addition, paint.rpd is not optimized for this release of Oracle Business Intelligence and should not be used as a basis for any repositories in production environments.

Using Online and Offline Repository Modes

You can open a repository for editing in either online or offline mode. The tasks you can perform depend on the mode in which you opened the repository.

To open a repository, you must enter the repository password. This password is specific to each repository and is used to encrypt the repository.

This section contains the following topics:

- [Opening a Repository in Offline Mode](#)
- [Opening a Repository in Online Mode](#)
- [Checking Out Objects](#)
- [Checking In Changes](#)
- [About Read-Only Mode](#)

Opening a Repository in Offline Mode

Use offline mode to view and modify a repository while it is not loaded into the Oracle BI Server. If you attempt to open a repository in offline mode while it is loaded into the Oracle BI Server, the repository opens in read-only mode. Only one Administration Tool session at a time can edit a repository in offline mode. See "[About Read-Only Mode](#)" for more information.

You do not need to enter a user name and password to open a repository in offline mode. You only need to enter the repository password.

Note: In offline editing, remember to save your repository from time to time. You can save a repository in offline mode even though the business models might be inconsistent.

To open a repository in offline mode:

1. In the Administration Tool, select **File > Open > Offline**.
2. Go to the repository you want to open, and then select **Open**.
3. In the Open Offline dialog, enter the repository password, and then click **OK**.

If the server is running and the repository you are trying to open is loaded, the repository opens in read-only mode. If you want to edit the repository while it is loaded, you must open it in online mode. Also, if you open a repository in offline

mode and then start the server, the repository becomes available to users. Any changes you make become available only when the server is restarted.

When you open a repository in the Administration Tool in offline mode, the titlebar displays the name of the open repository.

Opening a Repository in Online Mode

Use online mode to view and modify a repository while it is loaded into the Oracle BI Server. The Oracle BI Server must be running to open a repository in online mode. There are certain things you can do in online mode that you cannot do in offline mode. In online mode, you can perform the following tasks:

- Manage scheduled jobs
- Manage user sessions
- Manage the query cache
- Manage clustered servers

To open a repository in online mode:

1. In the Administration Tool, select **File > Open > Online** to display the Open Online Repository dialog.

The Oracle BI Server DSNs that have been configured on your computer are displayed in the dialog. If no additional DSNs have been configured for this version of the Oracle BI Server, you might see only the default DSN that is configured for you during installation.

See "Integrating Other Clients with Oracle Business Intelligence" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for information about how to create an ODBC DSN for the Oracle BI Server.

2. Enter the repository password for the repository currently loaded in the Oracle BI Server.

You can use Repository tab of the Deployment page in Fusion Middleware Control to view the name of the current repository.

3. Provide a valid user name and password.

The user you provide must have the ManageRepositories permission. See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for more information.

4. If you expect to work extensively with the repository (for example, you plan to check out many objects), select the **Load all objects on startup** option. This loads all objects immediately, rather than as selected. The initial connect time might increase slightly, but opening items in the tree and checking out items is faster.

5. Select the appropriate DSN and click **OK**.

When you open a repository in the Administration Tool in online mode, the titlebar displays the DSN for the Oracle BI Server to which you are connected, not the name of the current repository.

Guidelines for Using Online Mode

Use caution when making changes to your repository in online mode, and always perform a consistency check to ensure the repository is consistent before checking in objects.

In addition, you should limit the number of concurrent online users. The best practice is to have only one user working in online mode at a time. Even when users have different objects checked out, there might be dependencies between the objects that can cause conflicts when the changes are checked in. In general, only one user should make online changes in a single business model at a time.

If you must have multiple concurrent users in online mode, do not have more than five users. For situations where you need more than five users, use the multiuser development environment. See [Chapter 3, "Setting Up and Using the Multiuser Development Environment"](#) for more information.

Even with a single user making changes, be aware that online mode is riskier than offline mode because you are working against a running server. If you check in changes that are not consistent, it might cause the Oracle BI Server to shut down. When you work in online mode, make sure to have a backup of the latest repository so that you can revert to it if needed. You can also use **File > Undo All Changes** to roll back all changes made since the last check-in.

Checking Out Objects

When you are working in a repository open in online mode, you are prompted to check out objects when you attempt to perform various operations. Select the objects you want to check out and click **Yes** to check out the objects.

If you are performing a task in a wizard, the Checkout screen displays a summary of the objects that need to be checked out to complete the operation. Click **Next** to check out the objects and complete the task.

Checking In Changes

When you are working in a repository open in online mode, you are prompted to perform a consistency check before checking in the changes you make to a repository.

If you have made changes to a repository and then attempt to close the repository without first checking in your changes, a dialog opens automatically asking you to select an action to take. If you move an object from beneath its parent and then attempt to delete the parent, you are prompted to check in changes before the delete is allowed to proceed.

Use the Check in Changes dialog to make changes available immediately for use by other applications. Applications that query the Oracle BI Server after you have checked in the changes will recognize them immediately. Applications that are currently querying the server will recognize the changes the next time they access any items that have changed.

To make changes available and have them saved to disk immediately:

- In the Administration Tool, select **File**, then select **Check In Changes**.

If the Administration Tool detects an invalid change, a message is displayed to alert you to the nature of the problem. Correct the problem and perform the check-in again. Note that you can select a message row and click **Go To**, or double-click a message row, to go directly to the affected object.

You must save changes to persist the changes to disk. You must check in changes before you can save, but you do not need to save to check in changes.

About Read-Only Mode

Only one component (either the Oracle BI Server, or a single Administration Tool client in offline mode) can have a repository open in read/write mode at a time. If a second component opens a repository that is already in use, the repository is opened in read-only mode.

For example, assume the Oracle BI Server loads a repository in read/write mode. Any number of Administration Tool clients connecting to that repository in online mode will also get read/write mode, because they are accessing the repository through the Oracle BI Server. However, Administration Tool clients opening that repository in offline mode will get read-only mode, because the repository is already open for read/write through the Oracle BI Server.

Alternatively, assume an Administration Tool client opens a repository offline in read/write mode. When the Oracle BI Server starts, it will get read-only mode, as will any Administration Tool clients connecting to that repository in either offline or online modes. To enable the server to load the repository in read/write mode in this situation, you must first close the Administration Tool client that has the repository locked, and then restart the Oracle BI Server.

The Administration Tool also opens a repository in read-only mode when Oracle Business Intelligence has been clustered, and the Administration Tool is connected in online mode to a slave server. This occurs because the Master BI Server holds a lock on the repository. To avoid this situation when running in a clustered environment, ensure that the Oracle BI Server ODBC DSN used by the Administration Tool has been configured to point to the Cluster Controllers rather than to a particular Oracle BI Server.

In addition, the Administration Tool opens repositories in read-only mode when the configuration setting **Disallow RPD Updates** has been selected in Fusion Middleware Control. See "Using Fusion Middleware Control to Disallow RPD Updates" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about this configuration setting.

Checking the Consistency of a Repository or a Business Model

Repository metadata must pass a consistency check before you can make the repository available for queries. The Consistency Check Manager lets you enable and disable rules for consistency checks, find and fix inconsistent objects, and limit the consistency check to specific objects. You can also use the `validaterpd` utility to check the validity of all metadata objects.

This section contains the following topics:

- [About the Consistency Check Manager](#)
- [Checking the Consistency of Repository Objects](#)
- [Using the `validaterpd` Utility to Check Repository Consistency](#)

About the Consistency Check Manager

The Consistency Check Manager checks the validity of your repository to ensure that it can load at run time, and to identify any syntax or semantic errors that may cause queries to fail.

Each time you save the repository, a dialog asks if you want to check global consistency. You have the following options:

- **Yes.** Checks global consistency and then saves the repository file.
- **No.** Does not check global consistency and then saves the repository file.
- **Cancel.** Does not check global consistency and does not save the repository file.

The Consistency Check Manager does not check the validity of objects outside the metadata using the connection. It only checks the consistency of the metadata and not any mapping to the physical objects outside the metadata. If the connection is not working or objects have been deleted in the database, the Consistency Check Manager does not report these errors.

If you use lookup tables to store translated field names with multilingual schemas, note that consistency checking rules are relaxed for the lookup tables. See "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization and lookup tables.

The consistency checker returns the following types of messages:

- **Errors.** These messages describe errors that need to be fixed. Use the information in the message to correct the inconsistency, then run the consistency checker again. The following is an example of an error message:

```
[38082] Type of Hierarchy '"ORT_C41"..."ORT_C41/MDF_BW_Q02"."Product Hierarchy for Material MARA"' in Cube Table '"ORT_C41"..."ORT_C41/MDF_BW_Q02"' needs to be set.
```

If you disable an object and it is inconsistent, a message is displayed, asking if you want to make the object unavailable for queries.

- **Warnings.** These messages indicate conditions that may or may not be errors. For example, you might receive a warning message about a disabled join that was intentionally disabled to eliminate a circular join condition. Other messages may warn of inconsistent values, or feature table changes that do not match the defaults. The following is an example of a warning message:

```
[39024] Dimension '"Paint"."MarketDim"' has defined inconsistent values in its levels' property 'Number of elements'.
```

- **Best Practices.** These messages provide information about conditions, but do not indicate an inconsistency. The message appears if a condition violates a best practice recommendation.

Note: After upgrading from a previous software version and checking the consistency of your repository, you might notice messages that you had not received in previous consistency checks. This typically indicates inconsistencies that had been undetected before the upgrade, not new errors.

In the Consistency Check Manager, you can sort the rows of messages by clicking the column headings. Additionally, the status bar provides a summary of all the rows displayed.

Checking the Consistency of Repository Objects

You can use the Administration Tool to check consistency in the following ways:

- To check consistency for all objects in the repository, select **File**, then select **Check Global Consistency**.

- To check the consistency of a particular repository object, such as a physical database, business model, or subject area, right-click the object and select **Check Consistency**.
- If you already have the Consistency Check Manager open, you can check global consistency by clicking **Check All Objects**.

To view the Consistency Check Manager without performing a global consistency check, select **Tools**, then select **Show Consistency Checker**. If you have checked consistency in the current session, the messages from the last check appear in the Messages pane.

To check the consistency of a repository:

1. In the Administration Tool, select **File**, then select **Check Global Consistency**. The Consistency Check Manager is displayed, listing any messages relating to the current repository.

Note: If you disable an object and it is inconsistent, a dialog appears, asking whether you want to make the object unavailable for queries.

2. To edit the repository to correct inconsistencies, double-click any cell in a row to open the properties dialog for that object. Then, correct the inconsistency and click **OK**.
3. To copy the messages so that you can paste them in another file such as a spreadsheet, select one or more rows and click **Copy**. Note that clicking **Copy** without any rows selected copies all messages.
4. To check consistency again, click **Check All Objects** to perform a global check. Or click the **Refresh** button in the top right corner to check only the objects that were listed as inconsistent in the last check.
5. When finished, click **Close**.

To check the consistency of a single object in a repository:

1. In the Administration Tool, right-click an object, then select **Check Consistency**.
If the object is not consistent, a list of messages appears.
2. To edit the repository to correct inconsistencies, double-click any cell in a row to open the properties dialog for that object. Then, correct the inconsistency and click **OK**.
3. To copy the messages so that you can paste them in another file such as a spreadsheet, click **Copy**.
4. To check consistency of the object again, click the refresh button at the top right corner of the dialog.

If you click **Check All Objects**, all objects in the repository are checked.

Using the `validaterpd` Utility to Check Repository Consistency

You can use the Oracle BI Server utility `validaterpd` to check the validity of all metadata objects in a repository. You can run this utility on any platform supported by the Oracle BI Server. Running this utility performs the same validation checks as the Consistency Check Manager in the Administration Tool.

Before running `validaterpd`, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX systems) to launch a command prompt that is initialized to your Oracle instance. You can find this utility in:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup
```

Then, run `validaterpd` from the resulting command prompt with the desired options. You can run the utility from this directory with no arguments or parameters to see usage information.

The utility takes the following parameters:

```
validaterpd -R repository_name -O output_file_name [-P repository_password] [-8]
```

Where:

repository_name is the name and path of the repository that you want to validate.

output_file_name is the name and path of a text file where the validation results will be recorded.

repository_password is the password for the repository that you want to validate.

Note that the *repository_password* argument is optional. If you do not provide the password argument, you are prompted to enter the password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide password arguments either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.

`-8` specifies UTF-8 encoding in the output file.

Example

The following example generates an output file called `results.txt` that contains validation information for the repository called `repository.rpd`:

```
validaterpd -R repository.rpd -O results.txt  
Give password: my_rpd_password
```

Setting Up and Using the Multiuser Development Environment

Multiuser development (MUD) provides a mechanism for concurrent development on overlapping code bases. Oracle Business Intelligence provides a MUD environment that manages subsets of metadata, in addition to multiple users, by providing a built-in versioning system for repository development. This chapter explains how to set up and use the multiuser development environment in Oracle Business Intelligence, including defining projects, setting up the multiuser development directory, checking out and checking in projects, and merging metadata.

See also "[Managing the Repository Lifecycle in a Multiuser Development Environment](#)" for additional information about working in a multiuser development environment.

This chapter contains the following topics:

- [About the Multiuser Development Environment](#)
- [Setting Up Projects](#)
- [Setting Up the Multiuser Development Directory](#)
- [Making Changes in a Multiuser Development Environment](#)
- [Checking In Multiuser Development Repository Projects](#)
- [Branching in Multiuser Development](#)
- [Viewing and Deleting History for Multiuser Development](#)
- [Setting Multiuser Development Options](#)

About the Multiuser Development Environment

In Oracle Business Intelligence, multiuser development facilitates the development of application metadata in enterprise-scale deployments. Application metadata is stored in a centralized metadata repository (RPD) file. The Administration Tool is used to work with these repositories.

The following are examples of how you might use a multiuser development environment:

- Several developers work concurrently on subsets of the metadata and then merge these subsets back into a master repository without their work conflicting with other developers. For example, after completing an implementation of data warehousing at a company, an administrator might want to deploy Oracle Business Intelligence to other functional areas.

Note: In this chapter, "master repository" refers to the copy of a repository in the multiuser development directory

- A single developer manages all development. For simplicity and performance, this developer might want to use the multiuser development environment to maintain metadata code in smaller chunks instead of in a large repository.

In both examples, an administrator creates projects in the repository file in the Administration Tool, then copies this repository file to a shared network directory (called the multiuser development directory). Developers are able to check out projects, make changes and then merge the changes into the master repository. When developers check out projects using the Administration Tool, files are automatically copied and overwritten in the background. Therefore, it is important for the administrator to perform setup tasks and for the developers to perform check-out and check-in procedures carefully, paying close attention to the Administration Tool messages that appear.

When developers check out projects, repository files are not automatically copied or overwritten. Instead, the Administration Tool creates two new files when projects are checked out: one to hold the original project data, and one to hold the project changes.

For example, when a repository developer checks out project A from master.rpd in the C:\multiuser development directory, the Administration Tool extracts all metadata related to project A and prompts the developer for a new file name to save the data. When the developer chooses a new file name, for example Mychanges.rpd, the Administration Tool creates two new files:

- A file called MyChanges.rpd that will contain the changes made by the developer
- A file called originalMyChanges.rpd that contains the original project data

The originalMyChanges.rpd file helps determine what changes have been made by the developer in Mychanges.rpd. This information is needed during the multiuser development merge process.

Note: To reduce storage needs, repositories in Oracle Business Intelligence Enterprise Edition 11g Release 1 (11.1.1) are stored in a compressed format. Because of this, you may notice that the size of an RPD file opened and saved in this release is significantly smaller than the size of RPD files from previous releases.

About the Multiuser Development Process

Multiuser development presupposes a clear understanding of customer technical and business objectives. It also requires that you follow clearly defined development processes and adhere rigorously to those processes, including consistent merging and reconciliation practices.

The following procedure shows the general steps to follow when deploying a multiuser development environment. The first three steps are usually performed by an administrator, and the remaining steps are usually performed by one or more developers.

To deploy a multiuser development environment:

1. Define projects to organize voluminous metadata into manageable components. See "[Creating Projects](#)" for more information. Consider these tips:

- Use smaller RPDs to shorten and simplify development effort and unit testing.
- Organize development resources by projects to spread workload and reduce inconsistencies and overwrites.
2. Set up a shared network directory to use as the multiuser development directory.
3. Copy the master repository to the multiuser development directory.
4. Extract one or more projects for local development.
5. Merge repository objects and resolve conflicts.
 - Because metadata objects are often highly interrelated, several developers could be working on the same objects.
 - When configuration conflicts occur during check in, developers are prompted for the correct process.
 - The master repository is locked for check-ins when you merge your local changes. However, other developers can still perform check-outs.
6. Publish changes to the network.
 - Many developers can simultaneously work on the same objects, but only one can publish at a time. During the publish, no check-outs are allowed.
7. Use Logging and Backup features to identify points of erroneous or incorrect configuration.
 - The log file tracks multi-development activity, along with comments.
 - The master repository and developer repositories are automatically backed up for future reference and for use in manual rollback.

Setting Up Projects

Projects are the central enabler of metadata management. A project consists of a discretely-defined subset of the repository metadata, in the form of groups of logical stars with associated metadata. A project has the following characteristics:

- Is largely defined by logical fact tables in the applicable business model
- Automatically adds related logical dimension tables and other metadata during extract
- Can have one to many logical fact tables

For projects that are just beginning, the best practice is to begin with a repository containing all the necessary physical table and join definitions. In this repository, you create a logical fact table as a placeholder in the Business Model and Mapping layer and a subject area as a placeholder in the Presentation layer. As you add business model and subject area metadata, new projects based on individual subject areas and logical facts can be created.

Note: Only one person at a time can create projects in a master repository.

This section contains the following topics:

- [About Projects](#)
- [Creating Projects](#)

- [About Converting Older Projects During Repository Upgrade](#)

About Projects

Projects can consist of Presentation layer subject areas and their associated business model logical facts, dimensions, groups, users, variables, and initialization blocks. Administrators can create projects so that developers and groups of developers can work on projects in their area of responsibility.

The primary reason to create projects is to support multiuser development. During the development process, you can split up the work (metadata) between different teams within your company by extracting the metadata into projects so that each project group can access a different part of the metadata.

In addition to multiuser development, you may want to create projects for licensing reasons. Before releasing a new software version, you may want to ensure that only the metadata that is relevant to the licensed application is in a project and that everything is consistent and complete. You can accomplish this by adding only the fact tables that are relevant to the application.

Project extractions are fact table centric. This ensures that project extracts are consistent and makes licensing much easier to manage.

About the Project Dialog

In the Project dialog, the left pane contains objects that you can use to create a project. The objects in the right pane are all the objects you chose (directly or indirectly) that reflect the complete set of data that makes each addition consistent. For example, if you select a subject area from the top node of the left-hand tree to add to your project, underlying fact tables of other subject areas are automatically added if needed to make the extract consistent.

The following describes the left pane of the Project dialog:

- You can choose to group fact tables by Business Model or Subject Area, to help select the fact tables you want. Typically, grouping fact tables according to which ones are used by a particular subject area is a more convenient way to choose fact tables for your project. Note that a fact table can be associated with multiple subject areas, but belongs to one and only one business model.

Although it appears that you can add a subject area from the top node when you group facts by subject area, you are actually adding only the underlying fact tables. The subject areas only appear as choices to help you to add the elements you want in your project. Additionally, it adds any other objects that are necessary to make the extract consistent. To add an actual subject area, use the Presentation node at the bottom of the tree.

- When grouped by business model, the left pane displays only facts that belong to the business model.
- The Presentation node contains presentation layer objects. You must explicitly include these objects in your project if you want to work with them; they are not added automatically.

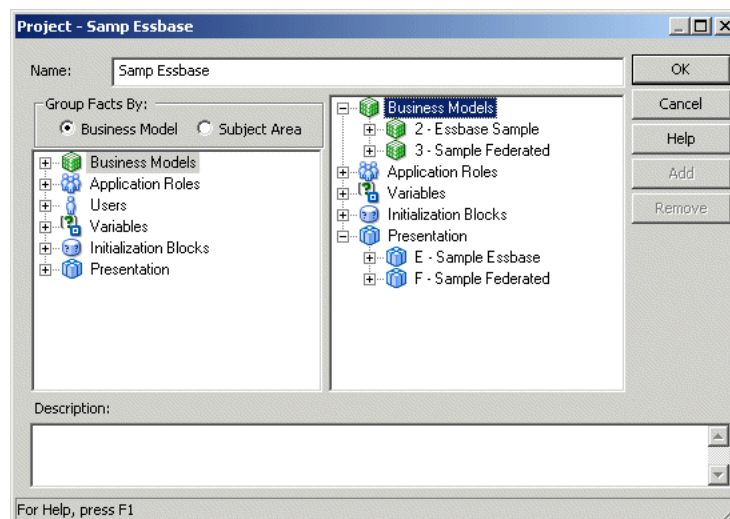
If you add presentation objects that are not related to any fact tables in the project, a warning appears when you click **OK**. The consistency checker also notes this discrepancy.

The right pane of the Project dialog shows the objects you select to be extracted, such as fact tables (under the Business Models folder), Presentation layer objects (under the

Presentation folder), users, application roles, variables, and initialization blocks. These objects are extracted when you click **OK**.

Figure 3–1 shows the Project dialog.

Figure 3–1 Project Dialog with Fact Tables Grouped by Business Model



Creating Projects

When creating a project, you typically select a subject area or a subset of logical fact tables related to the selected subject area, and the Administration Tool automatically adds any business model and Physical layer objects that are related. An object can be part of multiple projects. Alternatively, if you choose to group facts by business model, you can select a particular business model or a set of logical fact tables that are part of a business model. You also need to explicitly add any Presentation layer objects if you want them to be part of your project.

Note that although the project definition itself does not include Physical layer objects, these objects are extracted and determined through the project definition.

After you create projects, they become part of the metadata and are available to multiple developers who need to perform development tasks on the same master repository. When defined this way, projects typically become a consistent repository after a developer checks out the projects and saves them as a new repository file.

To create a project for a multiuser development environment:

1. In the Administration Tool, choose **File > Open > Offline**.
2. In the Open dialog, select the repository that you want to make available for multiuser development, then click **OK**. Provide the repository password, then click **OK** again.
3. Select **Manage**, then select **Projects**.
4. In the Project Manager dialog, in the right pane, right-click and then select **New Project**.

The left pane contains the objects that are available to be placed in a project. The right pane contains the objects that you select to be part of the project.

5. In the Project dialog, type a name for the project.

6. Choose whether to group facts by business model, or subject area. It is typically more convenient to group facts by subject area.
7. Perform one or more of the following steps to add fact tables to your project:
 - In the left pane, select a subject area or business model and then click **Add**. The Administration Tool automatically adds all the associated logical fact tables.
 - In the left pane, expand the subject areas or business models and select one or more logical fact tables that are related to the subject area or that are within the business model, then click **Add**.

The project is defined as explicitly containing the selected logical fact tables and implicitly containing all logical dimension tables that are joined to the selected logical fact tables (even though they do not appear in the right pane).

See "[About the Project Dialog](#)" for more information about the objects that appear in the left and right panes.

8. To remove fact tables from the project, in the right pane, select a fact table and click **Remove**. You can also remove all fact tables associated with a subject area or business model by selecting a subject area or business model and clicking **Remove**.
9. Optionally, add any application roles, users, variables, or initialization blocks needed for the project. Although objects like variables and initialization blocks that are directly referenced by other extracted objects are included automatically, you might want to include objects in your project that are not referenced. For example:
 - If you are using initialization blocks for authentication, include any necessary initialization blocks.
 - Include repository variables or other objects that are not yet referenced by other objects, but that you might want to use in future repository development.
 - Include users and application roles that are currently being used, or will be used in the future, as part of your data access security settings.

Tip: You may want to add the top node for each object type (for example, Variables), then selectively remove individual objects from the right pane.

10. Select the Presentation layer objects that you want to include in your project from the left pane and click **Add**. You must add these objects to see them in the project; they are not added automatically.

You can also remove particular presentation tables or columns from the project definition by double-clicking the object in the right pane, or selecting the object and clicking **Remove**.

Note: If you do not see the set of subject areas you expect after the project is created, edit the project to explicitly add the subject areas you need.

11. Click **OK**.

About Converting Older Projects During Repository Upgrade

When you upgrade a repository from Oracle Business Intelligence versions before 10.1.3.2, the project definition is upgraded. During the upgrade, the project definition, subject areas, target levels, list catalogs, and existing fact tables are automatically converted into simple fact tables in the following way:

- Get presentation columns related to the target levels through the qualifying keys.
- Get presentation columns related to the list catalogs through the qualifying keys.
- Get presentation columns related to the subject areas.
- Get all the logical columns from all the presentation columns.
- Get all the logical columns from the fact tables in the project.
- Get the fact tables from all the logical columns.

After the upgrade, projects contain only simple fact tables. All the security objects remain unchanged.

In addition, projects in repositories from any version before 11g Release 1 (11.1.1) are upgraded so that they explicitly contain Presentation layer objects. In previous releases, Presentation layer objects were implicitly included based on the permissions of the users included in the project.

Setting Up the Multiuser Development Directory

To prepare for multiuser development, an administrator performs the following tasks:

- Identify or create a shared network directory that will be dedicated to multiuser development.
- After creating all projects, copy the repository file in which you created the projects to the multiuser development directory where it will be used as your master repository for multiuser development.

After the administrator has identified the multiuser development directory and copied the repository file, developers must set up to the Administration Tool to point to the multiuser development directory before they can check out projects.

This section contains the following topics:

- [Identifying the Multiuser Development Directory](#)
- [Copying the Master Repository to the Multiuser Development Directory](#)
- [Setting Up a Pointer to the Multiuser Development Directory](#)

Identifying the Multiuser Development Directory

After defining all projects, the administrator must identify or create a shared network directory (called the multiuser development directory) that all developers can access, and then upload the new master repository to that location. This shared network directory should be used only for multiuser development. This directory typically contains copies of repositories that need to be maintained by multiple developers. The multiuser development directory must be on a Windows system.

Developers create a pointer to the multiuser development directory when they set up the Administration Tool on their computers.

Caution: The administrator must set up a separate, shared network directory that is dedicated to multiuser development. If not set up and used as specified, critical repository files can be unintentionally overwritten and repository data can be lost.

Copying the Master Repository to the Multiuser Development Directory

After the multiuser development directory is identified, the administrator must copy the master repository file to the multiuser development directory. Projects from this master repository will be extracted and downloaded by the developers who will make changes and then merge these changes back into the master repository.

After you copy the repository to the multiuser development network directory, notify developers that the multiuser development environment is ready.

Setting Up a Pointer to the Multiuser Development Directory

Before checking out projects, each developer must set up their Administration Tool to point to the multiuser development directory on the network. The Administration Tool stores this path in a hidden Windows registry setting on the workstation of the developer and uses it when the developer checks out and checks in objects in the multiuser development directory.

Note: Until the pointer is set up, the multiuser options are not available in the Administration Tool.

Initially, the network directory contains the master repositories. The repositories in this location are shared with other developers. Later, the network directory contains additional multiuser development history files, including historical subsets and repository versions. Do *not* manually delete any files in the multiuser development directory; these files are important and are used by the system.

When setting up the pointer, the developer can also complete the **Full Name** field. Although the field is optional, it is recommended that the developer complete this field to allow other developers to know who has locked the repository. The Full Name value is stored in `HKEY_CURRENT_USER` in the registry, and is unique for each login.

To set up a pointer to the multiuser development directory:

1. From the Administration Tool menu, choose **Tools > Options**.
2. In the Options dialog, click the Multiuser tab.
3. In the Multiuser tab, for **Multiuser development directory**, enter the full path to the network directory.

Alternatively, click **Browse**, select the multiuser development directory, and then click **OK**.

4. In the **Full Name** field, type your complete name, then click **OK**.

Making Changes in a Multiuser Development Environment

During check-out and check-in, a copy of the master repository is temporarily copied to the developer's local repository directory (typically, `ORACLE_INSTANCE\bifoundation\OracleBIServerComponent\coreapplication_obisn\repository` by default). After checking out projects and making changes in a

local repository file, each developer can check in (merge) changes into the master repository or discard the changes.

To make changes in a multiuser development environment, perform the tasks described in the following sections:

- [Checking Out Repository Projects](#)
- [About Changing and Testing Metadata](#)
- [About Multiuser Development Menu Options](#)

Checking Out Repository Projects

After setting up a pointer to the multiuser development default directory, a developer can check out projects, change metadata, and test the metadata. In the **File > Multiuser** submenu, the **Checkout** option is only available when there is a multiuser development directory defined in the Multiuser tab of the Options dialog.

If a developer checks out a local repository and attempts to exit the application before publishing it to the network or discarding local changes, a message appears to allow the developer to select an action. See "[About Closing a Repository Before Publishing It to the Network](#)" for more information.

This section contains the following topics:

- [About Repository Project Checkout](#)
- [Checking Out Projects](#)
- [Using the extractprojects Utility to Extract Projects](#)

About Repository Project Checkout

During checkout, the Administration Tool performs the following tasks:

- In the developer's local repository directory, the Administration Tool makes a temporary copy of the master repository.

Note: If a repository with that name exists in this location, the developer is asked to confirm overwriting the existing repository. If the developer clicks **Yes**, the existing local repository is immediately overwritten in the background and after the new repository is saved, the temporary master repository file is automatically deleted.

- In the developer's local repository directory, the Administration Tool saves a local copy of the selected projects in a new repository, such as Metadata1.rpd. The developer provides a name for the local copy. The developer makes metadata changes in this file. The number is incremented for each checkout for that session.
- In the developer's local repository directory, the Administration Tool saves a second local copy of the new repository, adding 'original' as the prefix (for example, originalMetadata1.rpd).
- After the developer saves the new repository file, check out is complete. In the developer's local repository directory, the temporary copy of the master repository is automatically deleted.

Caution: When the developer selects and saves the projects to a local repository file, the Administration Tool does not place a lock on the projects in the master repository on the shared network drive. Therefore, nothing physically prevents others from working on the same project. To determine if a project has been checked out, you need to look in the log file in the multiuser development directory on the shared network drive.

Checking Out Projects

This section explains how to check out projects using the Administration Tool.

To check out projects:


1. From the Administration Tool menu, choose **File > Multiuser > Checkout**.
2. If there is more than one repository in the multiuser development directory, the Multiuser Development Checkout dialog appears. Select the appropriate repository, and then click **OK**.

This dialog does not appear if there is only one repository in the multiuser development directory.

3. In the Extract from dialog, type the repository password, and then click **OK**.
If no projects exist in the repository, a message appears and the repository does not open.
4. If there is more than one project in the master repository, the Browse dialog appears. Select the projects that you want to be part of your project extract, and then click **OK**.

Figure 3–2 shows the Browse dialog for selecting projects.

Figure 3–2 Browse Dialog for Selecting Projects



If only one project exists in the master repository, it is selected automatically and the Browse dialog does not appear.

5. In the Create new subset repository dialog, type a name for the new repository (for example, Metadata1.rpd) and then click **Save**.

A working project extract repository is saved on your local computer. The name is exactly as you specified and is opened in offline mode. A log file is also created.

Caution: A second copy of the project extract repository is saved in the same location. The name of this version contains the word "original" added to the beginning of the name that you assigned to the repository extract. Do not change the original project extract repository. It is used during the multiuser development merge process, and when you want to compare your changes to the original projects.

Using the `extractprojects` Utility to Extract Projects

You can use the Oracle BI Server utility `extractprojects` to extract projects from a given repository without the overhead of the MUD environment. You can run this utility on any platform supported by the Oracle BI Server.

The `extractprojects` utility generates an RPD file that includes the set of projects you specify. The utility does not perform other tasks that are performed when you check out projects using the Administration Tool, like saving an original repository file or tracking the extract as a check-out in the MUD directory.

Before running `extractprojects`, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX systems) to launch a command prompt that is initialized to your Oracle instance. You can find this utility in:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup
```

Then, run `extractprojects` from the resulting command prompt with the desired options. You can run the utility from this directory with no arguments or parameters to see usage information.

The utility takes the following parameters:

```
extractprojects -B base_repository_name -O output_repository_name [-I input_  
project_name] [-P repository_password] [-L]
```

Where:

base_repository_name is the name and path of the repository from which you want to extract projects.

output_repository_name is the name and path of the repository generated by the utility.

input_project_name is the name of a project you want to extract. You can enter multiple projects. Be sure to precede each project entry with `-I` (for example, `-I project1 -I project2`). If the project name contains spaces, enclose it in double quotes (for example, "project 1").

repository_password is the password for the repository from which you want to extract projects.

Note that the *repository_password* argument is optional. If you do not provide the password argument, you are prompted to enter the password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide password arguments either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.

`-L` enables logging. When logging is enabled, a log file in the format `ProjExtr.YYYYMMDD.HHMMSS.xml` is created in the Oracle BI Server logging directory. For example:

ORACLE_INSTANCE/diagnostics/logs/OracleBIServerComponent/coreapplication_obisn/ProjExtr.20100616.082904.xml

Example

The following example extracts project1 and project2 from my_repos.rpd and creates a new repository called extract_repos.rpd:

```
extractprojects -B my_repos.rpd -O extract_repos.rpd -I project1 -I project2
Give password: my_rpd_password
```

Note: Be sure to provide the full pathnames to your repository files, both the input file and the output file, if they are located in a different directory.

About Changing and Testing Metadata

Most types of changes that can be made to standard repository files are also supported for local repository files. Developers can add new logical columns, logical tables, change table definitions, logical table sources, and so on. Developers may also work simultaneously on the same project locally. It is important to note, however, that Oracle Business Intelligence assumes the individual developer understands the implications these changes might have on the master repository. For example, if a developer deletes an object in a local repository, this change is propagated to the master repository when local changes are merged without a warning prompt.

To ensure metadata integrity, you should not remove a physical column unless there are no logical table source mappings to that physical column. Because of this, if you are using a multiuser development environment, you cannot delete a logical column and its associated physical column at the same time. Instead, you must first delete the logical column and perform a merge. Then, you can delete the physical column and perform another merge to safely remove the object.

You should not modify physical connection settings in a local repository. These are intentionally not propagated, and developers should not test the master connection pool settings in local environments. Instead, developers should apply settings for their local test data sources to perform unit testing of their model changes.

Physical connection settings, security settings, and database feature table changes are not retained in a multiuser development merge to prevent developers from overwriting passwords and other important objects in the master repository.

After making changes to a local repository, the developer can edit the local NQSConfig.INI file, enter the name of the repository as the default repository, and test the edited metadata.

Note: DSNs specified in the metadata must exist on the developer's workstation.

About Multiuser Development Menu Options

After the local developer makes changes, tests the changes, and saves the repository locally, the local developer can perform the following tasks from the **File > Multiuser** submenu:

- **Compare with Original.** Compares the working extracted local repository to the original extracted repository. When this option is selected, the Compare

repositories dialog opens and lists all the changes made to the working extracted repository since you checked out the projects.

- **Merge Local Changes.** Locks the master repository on the network multiuser directory to allow you to check in your changes. See "[Checking In Multiuser Development Repository Projects](#)" for more information.
- **Publish to Network.** After you successfully merge your changes, the master repository opens locally and the **Publish to Network** submenu item is available. When you select this option, the lock is removed, the repository is published, and the repository closes. See "[Checking In Multiuser Development Repository Projects](#)" for more information.
- **Undo Merge Local Changes.** Rolls back any previously merged local changes, and leaves the repository checked out so that you can make additional changes and then merge your local changes again. This option is only available after you have already merged local changes.
- **Discard Local Changes.** Any time after check out and before check in, you can discard your changes. When you select this option, the working repository closes without giving you an opportunity to save your work.

Caution: If you select this option, there is no opportunity to change your mind. For example, no confirmation dialog appears.

About Closing a Repository Before Publishing It to the Network

If you attempt to close an unpublished, locked repository without selecting an option in the **File > Multiuser** submenu, the Closing MUD repository dialog opens with the following options:

- **Publish repository.** Publishes the merged repository to the network share as the new master, releases the lock on the master, and the event is logged. This option is available after a Merge Local Changes event occurs. This option is also available on the **File > Multiuser** submenu.
- **Discard local changes.** Releases the lock on the master repository and records the event in the log. This option is available after a Checkout or Merge Local Changes is performed and can be found on the **File > Multiuser** submenu.
- **Close repository and keep lock.** This closes the repository, leaving the master repository locked.
- **Undo merge local changes.** Rolls back your previously merged local changes, and leaves the repository checked out so that you can make additional changes and then merge your local changes again.

Checking In Multiuser Development Repository Projects

After changing and testing the metadata on a local computer, the developer must check the projects into the master repository in the multiuser development directory. Only one developer at a time can merge metadata from a local repository into the master repository. Therefore, the master repository is locked at the beginning of the merge process.

The Oracle BI repository development process uses a three-way merge to manage concurrent development. Metadata merges are done first on local environments and then merged with the master repository. A three-way merge identifies local changes based on the following repository characteristics:

- The Master RPD
- The Baseline RPD or Master RPD snapshot at time of project extraction
- The current locally developed and changed RPD

Changes are managed by merge and reconciliation. Most of the merging process is automatic, and changes do not conflict. In case of any conflicting metadata sources, developers can locate and resolve them.

An administrator can also merge the changes from multiple repositories manually, or import objects from different repositories outside of a particular MUD environment.

Make sure to merge your changes frequently. The merge process is very complex and can become difficult if there are too many changes. See [Appendix D, "Merge Rules"](#) for more information about how objects are merged during the merge process.

This section contains the following topics:

- [About the Multiuser Development Merge Process](#)
- [Checking In Projects](#)
- [Tracking Changes to the Master Repository](#)

About the Multiuser Development Merge Process

The merge process involves the following files:

- **Original of the local (subset) repository.** Contains the state of the projects as originally extracted. This repository name begins with "original." An example of the file name for this copy might be originalDevelopment2.rpd. This version is stored in the same location as the modified (or working) version of the local repository.
- **Modified local (subset) repository.** Contains the extracted projects after being modified by the developer. This version is stored in the same location as the original version of the local repository.
- **Latest master repository in the multiuser development directory.** Note that this file may have been modified by other developers before this merge.

During the merge, the Administration Tool checks for added objects and if found, a warning message appears. The following list describes what happens during this step:

- **Warning about added objects.** When a person checks out a project, they have the ability to modify that project in any way and check it back in. Deletions and modifications are ways in which the integrity of the project is maintained. However, adding objects might introduce objects into the repository that do not belong to any project. Therefore, all project related objects are checked and if a new object is found, a warning message appears.

Caution: You must add newly created metadata to the project definition in the master repository for it to be visible in future extracted versions of the project. For example, if a developer checks out a project, adds a new object, and then checks it in, the new object is not visible in extracted versions of the project until it is explicitly added to the project definition. See "[Creating Projects](#)" for instructions.

- **Aggregation of related objects.** In the warning message, only the parent object is reported. The Administration Tool aggregates all the objects to make the message

more usable. For example, if a developer added a new business model, only the business model appears in the warning message to the user, not the tables, columns, dimensions, and so on.

When a developer publishes changes to the network, the following actions occur:

- The master repository in the multiuser development directory is overwritten with the repository containing the developer's changes.
- The *master_repository.lck* file is deleted. If another developer checks out the changed project from the master repository, the changes made by the first developer are exposed to the other developer.

How are Multiuser Merges Different from Standard Repository Merges?

The multiuser development check-in process uses the same technology as the standard repository merge process with a few important differences. See "[Performing Full Repository Merges](#)" for more information about the standard repository merge.

The following list describes the differences that occur during a multiuser development merge:

- Inserts (created objects) are applied automatically. Because a subset of the master repository is being used as the original repository, most objects in the master repository appear to be new. This would result in many unnecessary prompts that the developer would have to manually approve. Therefore, new objects are created without a prompt during a multiuser development merge.
- Conflicts that are not inserts but are resolved because of the automatic inserts are applied without a prompt during a multiuser development merge.
- The database and connection pool properties in the master repository take precedence over the same properties on the developer's computer. This precedence are applied without a prompt during a multiuser development merge.
- Changes to security settings are not retained when you perform a MUD merge to prevent developers from overwriting passwords and other important objects in the master repository.

To change security settings or database features in a multiuser development environment, you must edit the master repository directly. To do this, remove the master repository from the multiuser development directory, edit it in offline mode, then move it back.

Checking In Projects

When the check-in process begins, the following actions occur:

- The Administration Tool determines if the master repository is currently locked. If not, it locks the master repository, preventing other developers from performing a merge until the current merge is complete, and records the lock in the log file.
- For other developers, the **Merge Local Changes** option on the **File > Multiuser** menu is unavailable until the current check-in process has been successfully completed.
- The Administration Tool automatically copies the current version of the master repository from the multiuser development directory to the local repository directory on the developer's computer (typically *ORACLE_BI_HOME\oraInst\bifoundation\OracleBIServerComponent\coreapplication\repository*) and updates the log files in the local and multiuser development directories.

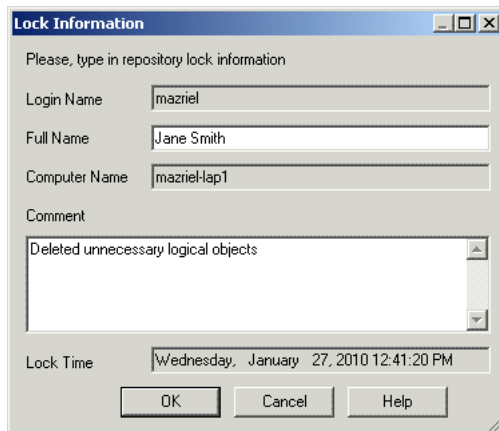
This is necessary because the master repository in the multiuser development directory might have changed after the developer checked out the projects.

To check in projects to the master repository:

1. In the Administration Tool, select **File > Multiuser > Merge Local Changes**, then click **Yes** if prompted to save changes.
2. In the Lock Information dialog, in the **Comment** field, type a description of the changes that you made, then click **OK**.

Figure 3–3 shows the Lock Information dialog.

Figure 3–3 Lock Information Dialog



3. If there are any conflicts, the Merge Repository Wizard opens and displays the Define Merge Strategy screen. Make merge decisions about whether to include or exclude objects by choosing **Current** or **Modified** from the **Decision** list. When you select an object in the decision table, the read-only text box below the decision table describes what changes were made to that object in the current repository. You can also click **View Change Statistics** to see a summary of changes. Click **Finish** when you are finished making merge decisions.

See "[Performing Full Repository Merges](#)" for additional information about the Define Merge Strategy screen.

A lack of conflicts does not mean that there are no differences between the repositories. Instead, it means that there are no decisions that have to be explicitly made by the developer to check in changes. See "[How are Multiuser Merges Different from Standard Repository Merges?](#)" for information about conflicts that are resolved automatically in a MUD merge.

In both cases, a CSV file is created in the local directory that contains details of the merged changes.

4. After you confirm all the changes, click **Save**.

This saves the merged repository locally, and then uploads this repository to the multiuser development directory with an incremented numeric file extension (for example, `Master_Sales.000`, `Master_Sales.001`, and so on).

At this point, the changes made by the developer are still not saved to the master repository in the multiuser development directory.

5. To commit these changes to the master repository in the multiuser development directory, select **File > Multiuser > Publish to Network**, then click **OK**.

The master repository in the multiuser development directory is overwritten with the copy of the repository containing the developer's changes.

Tracking Changes to the Master Repository

You can find a summary of the development activities on the master repository in *master_repository.log*. This log contains a record of the following activities:

- Projects that have been checked in and checked out and when these actions occurred
- The Windows login name and computer name initiating the transaction
- When locks are created and removed

Branching in Multiuser Development

Branching is a further refinement of the merging development process. Branching can provide higher efficiencies over large development teams that have overlapping releases, but it requires significant administrative overhead.

This section contains the following topics:

- [About Branching](#)
- [Using the Multi-Team, Multi-Release Model in Oracle Business Intelligence](#)
- [Synchronizing RPD Branches](#)

About Branching


In branching, developers work on private branches to isolate their code from other developers and merge changes back to the main branch. Different strategies can be followed, depending on the size of the development team.

In the **Simple Development Model**, all development occurs on a single main branch. This strategy has the following characteristics:

- Only for emergency fixes
- Checkouts may not be most current code
- Carries a stability risk for the mainline branch

Figure 3-4 shows the Simple Development Model.

Figure 3-4 Simple Development Model




In the **Small Team Development Model**, development occurs on a single Dev branch, with a separate Main branch strictly for releases. This strategy has the following characteristics:

- The Mainline is the official release branch
- Development occurs on a separate branch
- Stable code is merged back to Main at key milestones
- Branches are synchronized periodically

Figure 3-5 shows the Small Team Development Model.

Figure 3-5 Small Team Development Model

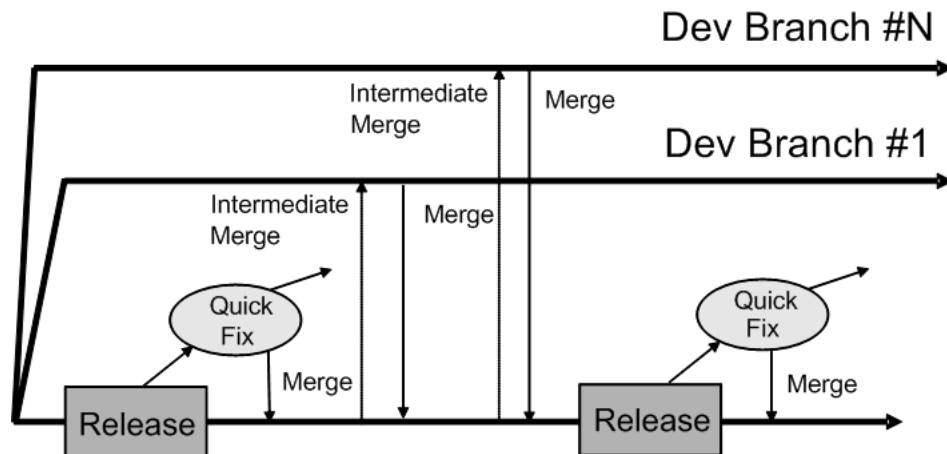


In the **Multi-Team, Multi-Release Model**, development occurs on multiple Dev branches, again with a separate Main branch strictly for releases. This strategy has the following characteristics:

- Supports more efficiency over disparate teams
- Development occurs on separate branches
- Stable code is merged back to Main at key milestones
- Branches are synchronized periodically

Figure 3-6 shows the Multi-Team, Multi-Release Model.

Figure 3-6 Multi-Team, Multi-Release Model



Using the Multi-Team, Multi-Release Model in Oracle Business Intelligence

Using complex branching strategies in Oracle Business Intelligence requires attentive organization of repository files, as well as altering the Multiuser setting in the Administration Tool. The following procedure provides an overview of the required steps.

To use the multi-team, multi-release model branching strategy:

1. Create a Main repository (Master Repository) and store it in the Master multiuser development directory.
 - Projects must be explicitly defined.
 - Branch developers should not have access to the Master directory.
2. Create a subset of branch repositories by extracting from Main and storing them as the Team1 and Team2 multiuser development directories. The Main and Team RPDs must be stored and secured in separate directories on the network.
3. Developers must check out, develop, merge, and publish from their respective Team RPDs. Developers A1 through A3 and B1 through B3 should manage their metadata work and merge to their Team repository.
 - Teams 1 and 2 must maintain their own repositories and periodically synchronize from Main to Team branches.
 - The Team repositories must be merged back into and published in the Main repository.
4. One specific group (for example, release management) should manage all project definitions, perform merges, publish, and synchronize the Team RPDs back to Main.

Synchronizing RPD Branches

For large development teams, it is a good practice to perform periodic branch synchronization as Main changes, in order to ease the ultimate Team check-in. Use the Administration Tool to synchronize repositories in a three-way merge.

To synchronize repository branches:

1. Check in all changes from your Team development branch and open the RPD in the Administration Tool. This is the **current** repository.
2. Extract a fresh Branch subset from Main. This is the **modified** repository.
3. In the Administration Tool, select **File**, then select **Merge** and browse to the backup of the previous Branch subset. This is the **original** repository.
4. Resolve all issues and perform the merge.

The RPD named in the **Save merged repository as** field becomes the new branch development RPD and is called the Original in future synchronizations.

Viewing and Deleting History for Multiuser Development

You can view and delete the development history of a multiuser development repository.

This section contains the following topics:

- [Viewing Multiuser Development History](#)

- [Deleting Multiuser Development History](#)

Viewing Multiuser Development History

You can view the development history of a multiuser development repository. In the Administration Tool, multiuser development history is only available when no repository is open and after the administrator sets up the shared network directory. This prevents the confusion that could occur if a user opened a history log that did not match an open, unrelated repository.

To view multiuser development history:

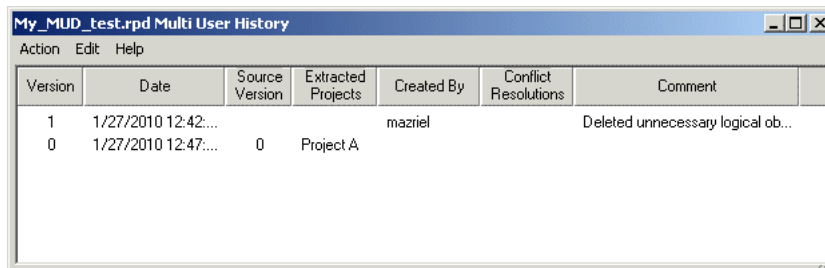
1. Open the Administration Tool.
2. Without opening a repository, select **File > Multiuser > History**.
3. In the Multiuser Development History dialog, select a repository.

A list of all master repositories in the multiuser development directory appears. If directory contains only one master repository, it is selected by default, and no list appears.

4. In the Open Offline dialog, type the password for the repository. The Multi User History dialog appears.

[Figure 3-7](#) shows the Multi User History dialog.

Figure 3-7 Multi User History Dialog



5. In the Multi User History dialog, right-click a row and select an option. [Table 3-1](#) describes the options in the Multi User History dialog.

Tip: To see details for all revisions, right-click in the background with no rows selected and select **View > Details**.

Table 3-1 Options in the Multi User History Dialog

Action	Description
View > Repository	Loads the selected master version of the repository in the Administration Tool in read-only mode.
View > Prior to Merge > Projects	Loads the selected version of a modified subset repository in the Administration Tool in read-only mode.
View > Conflict Resolution	Loads all necessary repositories of the selected version. Also shows the Merge dialog in read-only mode with all selected decisions as they were during the Merge Local Changes activity at that time. Double-clicking a row for a version with conflict resolutions has the same effect as selecting this menu item. Note: This menu item is only enabled for versions that had conflict resolutions.

Table 3–1 (Cont.) Options in the Multi User History Dialog

Action	Description
View > Details	Displays a log with details for the selected versions, or displays details for all versions if no specific versions are selected.
View > Prior to Merge > Changes	Compares modified subset repository of the selected version with original subset repository and shows all changes made by the user in the selected version.
Find and Find Again	Lets you search the list.
Select All	Selects all items displayed in the dialog.
Delete	Available only to multiuser development administrators.

Deleting Multiuser Development History

Only multiuser development administrators can delete history. Administrators are defined in a special hidden option file in the multiuser development directory. See ["Setting Multiuser Development Options"](#) for more information.

An administrator can delete the entire MUD history, or the oldest 1 to n versions. It is not possible to delete versions in the middle of the range. For example, an administrator cannot delete version 3 if there are still versions 1 and 2. If an administrator deletes the entire MUD history, newly assigned version numbers restart at version 1.

Setting Multiuser Development Options

You can create a multiuser development option file to specify options for multiuser development. The option file is a text file, in standard Windows INI format. It has the following properties and characteristics:

- The option file must be placed in the multiuser development directory. The file has the same name as the corresponding master repository, but with an .opt extension. For example, for `\\network\MUD\sales.rpd`, create an option file called `\\network\MUD\sales.opt`.
- The file should have the **Hidden** flag turned on.
- In general, the option file should be managed only by multiuser development administrators. To ensure this, you may want to change the sharing permissions for the file.

The following example shows a multiuser development option file:

```
[Options]
BuildNumber = Yes
Admin = admin1;admin2
Mandatory Consistency Check = Yes
Equalize During Merge = Yes
```

Options that are not explicitly set are turned off by default. To turn an option on, set its value to Yes. To turn an option off, either remove it from the option file, or set its value to No.

[Table 3–2](#) explains the options in the multiuser development option file.

Table 3–2 Options in the Multiuser Development Option File

Option	Description
BuildNumber	When set to Yes, the build version of the Administration Tool is displayed in the MUD history.
Admin	Lists multiuser development administrators. Administrators must be defined in the option file before they can delete MUD history. Administrators are defined by their computer/network login names. When multiple administrators exist, administrator names are separated by semicolons. For example: Admin=jsmith;mr Ramirez;plafleur
Mandatory Consistency Check	When set to Yes, the publish step performs a consistency check. Publishing cannot proceed unless there are no errors in the given repository.
Equalize During Merge	When set to Yes, the multiuser development merge process performs mandatory equalization during MUD merges. Note that setting this option to Yes affects the performance of the merge process.

Importing Metadata and Working with Data Sources

If you do not already have an Oracle BI repository file, you must create one. Then, you can import metadata from your data sources into the Physical layer of your repository. The Physical layer of the Administration Tool defines the data sources to which the Oracle BI Server submits queries, and the relationships between physical databases and other data sources that are used to process multiple data source queries.

Metadata imports to an Oracle BI repository must occur through an ODBC or native database connection to the underlying data source. Metadata can also be imported from software such as Microsoft Excel through an ODBC connection.

Importing metadata directly from each data source saves you time and effort by importing the structure for the Physical layer. Data from these sources can be displayed on Oracle BI Interactive Dashboards and other clients. You can only import metadata from supported data sources.

After you import metadata, properties in the associated database object and connection pool are typically set automatically. However, you may want to adjust database or connection pool settings. See [Chapter 6, "Setting Up Database Objects and Connection Pools"](#) for more information.

Although you can create the Physical layer manually rather than importing metadata, it is a labor-intensive and error-prone activity. It is strongly recommended that you import metadata.

This chapter contains the following topics:

- [Creating New Oracle BI Repository Files](#)
- [Performing Data Source Preconfiguration Tasks](#)
- [Importing Metadata from Relational Data Sources](#)
- [Importing Metadata from Multidimensional Data Sources](#)
- [Importing Metadata from XML Data Sources](#)
- [Using a Standby Database with Oracle Business Intelligence](#)

Creating New Oracle BI Repository Files

You can use the Create New Repository Wizard in the Administration Tool to create new Oracle BI repository files (.rpds). You do not need to create a new repository file if you already have an existing repository.

If you have an existing repository file, you can use the existing data source settings in that file as a template to connect to different data sources. To do this, use the existing data source settings and just change the database type and connection pool information. See ["Setting Up Database Objects"](#) and ["Creating or Changing Connection Pools"](#) for details.

To create a new repository file:

1. In the Administration Tool, select **File**, then select **New Repository**. The Create New Repository Wizard appears.

If an existing repository is open, you are prompted to save your changes, and the existing repository is closed.
2. Type a name for the repository. Keep the name to 156 characters or less to avoid problems with the metadata dictionary URL. An RPD file extension is automatically added if you do not explicitly specify it.
3. Select a location for the new repository. By default, new repositories are stored in the repository subdirectory, located at `ORACLE_INSTANCE\bifoundation\OracleBIServerComponent\coreapplication_obisn\repository`.
4. If you want to import metadata into the repository now, select **Yes** (the default) for **Import Metadata**. If you do not want to import metadata, select **No**.
5. Enter and confirm the password you want to use for this repository. The repository password must be longer than five characters and cannot be empty.

You enter the repository password when you open the repository in online or offline mode. It is used to encrypt the repository contents.
6. If you selected **Yes** for **Import Metadata**, click **Next**.

Refer to the following sections for information about the Import screens, according to your data source type:
 - [Importing Metadata from Relational Data Sources](#)
 - [Importing Metadata from Multidimensional Data Sources](#)
 - [Importing Metadata from XML Data Sources](#)
 - [Working with ADF Business Component Data Sources](#)Note that you may need to set up your data sources before you import information into the repository. See ["Performing Data Source Preconfiguration Tasks"](#) for more information.
7. If you selected **No** for **Import Metadata**, click **Finish** to create an empty repository.

Performing Data Source Preconfiguration Tasks

Before you can import physical objects from your data sources into your repository file, or set up connection pools to your data sources, you might need to perform configuration steps so that the data sources can be accessed by Oracle Business Intelligence.

For many data sources, you need to install client components. Client components are typically installed on the computer hosting the Oracle BI Server for query access, and on the computer hosting the Administration Tool (if different) for offline operations

such as import. In some cases, client components must be installed on the computer where the JavaHost is located.

Note: See also the following related topics:

- If the Oracle BI Server is running on a non-Windows platform, see [Chapter 15, "Setting Up Data Sources on Linux and UNIX"](#) for additional instructions.
 - See [Chapter 5, "Working with ADF Business Component Data Sources"](#) for information about setting up ADF Business Component data sources.
 - See ["System Requirements and Certification"](#) for information about the data source versions supported by Oracle Business Intelligence.
-
-

This section contains the following topics:

- [Setting Up ODBC Data Source Names \(DSNs\)](#)
- [Setting Up Oracle Database Data Sources](#)
- [Setting Up Oracle OLAP Data Sources](#)
- [Setting Up Oracle TimesTen In-Memory Database Data Sources](#)
- [Setting Up Essbase Data Sources](#)
- [Setting Up Hyperion Financial Management Data Sources](#)
- [Setting Up SAP/BW Data Sources](#)
- [Setting Up Oracle RPAS Data Sources](#)

Setting Up ODBC Data Source Names (DSNs)

Before you can import from a data source through an ODBC connection, or set up a connection pool to an ODBC data source, you must first create an ODBC Data Source Name (DSN) for that data source on the client computer. You reference this DSN in the Import Metadata Wizard when you import metadata from the data source.

You can only use ODBC DSNs for import on Windows systems.

To set up an ODBC DSN on Windows:

1. Open the Windows Control Panel by selecting **Start > Settings > Control Panel**, double-click **Administrative Tools**, and then double-click **Data Sources (ODBC)**.
2. In the ODBC Data Source Administrator dialog, click the System DSN tab, and then click **Add**.
3. From the Create New Data Source dialog, select the driver appropriate for your data source, and then click **Finish**.

The remaining configuration steps are specific to the data source you want to configure. Refer to the documentation for your data source for more information.

ODBC DSNs on Windows systems are used for both initial import, and for access to the data source during query execution. On UNIX systems, ODBC DSNs are only used for data access. For information about setting up ODBC data sources on UNIX, see [Chapter 15, "Setting Up Data Sources on Linux and UNIX."](#)

Note: For Teradata data sources, after you have installed the latest Teradata ODBC driver and set up an ODBC DSN, you must add the lib directory for your Teradata data source to your Windows system Path environment variable. For example:

D:\Program Files\NCR\Common Files\Shared ICU Libraries for Teradata\lib

Setting Up Oracle Database Data Sources

You should always use OCI when importing metadata from or connecting to an Oracle Database. Before you can import schemas or set up a connection pool, you must add a TNS names entry to your tnsnames.ora file. See the Oracle Database documentation for more information.

When you import metadata from an Oracle Database data source or set up a connection pool, you can include the entire connect string for **Data Source Name**, or you can use the net service name defined in the tnsnames.ora file. If you choose to enter only the net service name, you must set up a tnsnames.ora file in the following location within the Oracle Business Intelligence environment, so that the Oracle BI Server can locate the entry:

`ORACLE_HOME\network\admin`

Setting Up Oracle OLAP Data Sources

Before you import from an Oracle OLAP data source, ensure that the data source is a standard form Analytic Workspace.

In addition, the JavaHost process must be running to import from Oracle OLAP data sources, for both offline and online imports.

Setting Up Oracle TimesTen In-Memory Database Data Sources

Oracle TimesTen In-Memory Database is a high-performance, in-memory data manager that supports both ODBC and JDBC interfaces. These preconfiguration instructions assume that you have already installed TimesTen; see *Oracle TimesTen In-Memory Database Installation Guide* for more information.

To set up TimesTen data sources:

1. On the computer where TimesTen has been installed, create a Data Manager DSN (as a system DSN). See "Defining a Data Manager DSN" in *Oracle TimesTen In-Memory Database Operations Guide* for more information.
2. Perform an initial connection to the data store to load the TimesTen database into memory, and then create users and grant privileges, if you have not done so already. See "Managing Access Control" in *Oracle TimesTen In-Memory Database Operations Guide* for more information. Note that the default user of the data store is the instance administrator, or in other words, the operating system user who installed the database.
3. On the computer running the Oracle BI Server, install the TimesTen Client. See *Oracle TimesTen In-Memory Database Installation Guide* for more information.
4. On the computer where the TimesTen Client has been installed, create a Client DSN (as a system DSN). See "Creating Client DSNs" in *Oracle TimesTen In-Memory Database Operations Guide* for more information.

See "[Configuring Database Connections Using Native ODBC Drivers](#)" for information about how to perform this step when the Oracle BI Server is running on Linux or UNIX.

Note that if the TimesTen database is installed on the same computer as the TimesTen client, you can specify either the Data Manager DSN or the Client DSN in the Import Metadata Wizard.

Avoiding Query Failures Due to Limited Connections with TimesTen Data Sources

In some cases, queries to TimesTen data sources might fail with errors similar to the following:

```
[10058][State: S1000] [NQODBC] [SQL_STATE: S1000] [nQSError: 10058] A general
error has occurred.
[nQSError: 43113] Message returned from OBIS.
[nQSError: 43119] Query Failed:
[nQSError: 16023] The ODBC function has returned an error. The database may not be
available, or the network may be down.
Statement execute failed
```

To avoid these failures, increase the maximum number of connections for the TimesTen server, as follows:

1. In your TimesTen environment, open the `ttendaemon.options` file for editing. You can find this file at:

```
install_dir/srv/info
```

2. Add the following line:

```
-MaxConnsPerServer number_of_connections
```

To determine *number_of_connections*, use the following formula: if there are M connections for each connection pool in the RPD, N connection pools in the RPD, and P Oracle BI Servers, then the total number of connections required is $M * N * P$.

3. Save and close the file.
4. In the ODBC DSN you are using to connect to the TimesTen server, set the Connections parameter to the same value you entered in Step 2:
 - On Windows, open the TimesTen ODBC Setup wizard from the Windows ODBC Data Source Administrator. The Connections parameter is located in the First Connection tab.
 - On UNIX, open the `odbc.INI` file and add the Connections attribute to the TimesTen DSN entry, as follows:

```
Connections=number_of_connections
```

5. Stop all processes connecting to TimesTen, such as the `ttisql` process and the Oracle BI Server.
6. Stop the TimesTen process.
7. After you have verified that the TimesTen process has been stopped, restart the TimesTen process.

Setting Up Essbase Data Sources

The Oracle BI Server uses the Essbase client libraries to connect to Essbase data sources. You must ensure that the Essbase client libraries are installed on the computer running the Oracle BI Server before you can set up a connection to or import metadata from Essbase data sources. You also need to ensure that the Essbase client libraries are installed on any computer where you want to run the Administration Tool or the `nqcmd` utility. See ["System Requirements and Certification"](#) for information about supported versions of the Essbase Client for use with Oracle Business Intelligence.

After you verify that the Essbase client libraries are installed on the appropriate computers, you must ensure that the `PATH` environment variable on each computer includes the location of the Essbase client driver (for example, `EPM_ORACLE_HOME/products/Essbase/EssbaseClient`). See ["Configuring Essbase Data Sources on Linux and UNIX"](#) for information about setting the `PATH` variable on Linux.

You also need to ensure that an additional environment variable is set appropriately for each computer (either `ESSBASEPATH` or `ARBORPATH`, depending on your client version). For more information, see *Oracle Hyperion Enterprise Performance Management System Installation and Configuration Guide* (or the equivalent title for your client version).

Other system processes need the information contained in the environment variables set for the Essbase client libraries. The following sections provide additional information:

- [Updating Essbase Information in opmn.xml](#)
- [Adding Essbase Information to bi-init.cmd](#)

Updating Essbase Information in opmn.xml

To enable connectivity, on each computer running the Oracle BI Server, you need to verify that information in `opmn.xml` is correct for your Essbase client version and installed location.

To include the location of the Essbase client libraries in opmn.xml:

1. Open `opmn.xml` for editing. You can find `opmn.xml` at:

```
ORACLE_INSTANCE/config/OPMN/opmn/opmn.xml
```

2. Locate the `ias-component` tag for the Oracle BI Server process. For example:

```
<ias-component id="coreapplication_obis1" inherit-environment="true">
```

3. Under the Oracle BI Server `ias-component` tag, locate the line that identifies the variable `ESSBASEPATH`. For example:

```
<variable id="ESSBASEPATH" value="$ORACLE_HOME/products/Essbase/EssbaseServer" />
```

4. If necessary, update the value to point to the appropriate location for your installed client libraries. Change the variable name to `ARBORPATH` if necessary, depending on your client version.

5. Locate the line that identifies the variable `PATH`. For example:

```
<variable id="PATH" value="$ORACLE_HOME/bifoundation/server/bin:$ORACLE_HOME/bifoundation/web/bin:$ORACLE_HOME/bifoundation/odbc/lib:$ORACLE_INSTANCE:$ORACLE_HOME/lib:$ESSBASEPATH/bin:$PATH:$bin" append="true" />
```

6. Locate the portion of the entry that shows the location of the Essbase client libraries (shown in bold in the preceding example), and update the location if appropriate for your installation. Change the variable used in the path from ESSBASEPATH to ARBORPATH if necessary, depending on your client version.
7. Save and close the file.
8. Restart OPMN.
9. Repeat these steps on each computer that runs the Oracle BI Server process. If you are running multiple Oracle BI Server instances on the same computer, be sure to update the ias-component tag appropriately for each instance in opmn.xml (for example, ias-component id="coreapplication_obis1", ias-component id="coreapplication_obis2", and so on).

Adding Essbase Information to bi-init.cmd

On Windows systems, to enable connectivity, you need to update bi-init.cmd on each computer running the Administration Tool.

See "[Configuring Essbase Data Sources on Linux and UNIX](#)" for information about setting up the environment on Linux systems.

To add Essbase environment variables to bi-init.cmd:

1. Open the bi-init.cmd file for editing. You can find bi-init.cmd at:

```
ORACLE_INSTANCE\bifoundation\OracleBIApplication\coreapplication\setup
```

2. Add a line setting the ESSBASEPATH or ARBORPATH environment variable, as appropriate for your version and installed location of the Essbase client libraries. For example:

```
set ORACLE_INSTANCE=C:\mw_home\instances\instance1
set ORACLE_BI_APPLICATION=coreapplication
set ESSBASEPATH=C:\products\Essbase\EssbaseClient
call %ORACLE_INSTANCE%\bifoundation\OracleBIApplication\%ORACLE_BI_
APPLICATION%\setup\user.cmd
```

3. Locate the line that sets the PATH environment variable, and add the location of the Essbase client binaries. For example:

```
set PATH=%ORACLE_HOME%\bifoundation\server\bin;%ORACLE_
HOME%\bifoundation\web\bin;%ORACLE_HOME%\bin;%ESSBASEPATH%\bin;%PATH%
```

4. Save and close the file.

Setting Up Hyperion Financial Management Data Sources

Before you can import from or set up a connection to Hyperion Financial Management data sources, you must ensure that the Hyperion Financial Management win32 client is installed on the computer running the Oracle BI JavaHost process. If you are using Hyperion Financial Management version 9.3.1, you must also install the HFM Application Builder version 9.2.0.1 component on the JavaHost computer.

When you install the Hyperion Financial Management win32 client, make sure to perform the following steps in the Server/Cluster Registration tab of the client configuration utility:

- Provide the IP address or host name of the Hyperion Financial Management server.
- Click **Enable DCOM**.

In addition, you must edit the file `opmn.xml` on each computer running the Oracle BI JavaHost process to include environment variables required by Hyperion Financial Management. Note that the JavaHost process must be running to import from Hyperion Financial Management data sources, for both offline and online imports.

Note: You can only use Hyperion Financial Management data sources with Oracle BI EE when Oracle BI EE is running in a Windows-only deployment.

For Hyperion Financial Management 9.3.1:

1. Open `opmn.xml` for editing. You can find `opmn.xml` at:

```
ORACLE_INSTANCE/config/OPMN/opmn/opmn.xml
```

2. Locate the `ias-component` tag for the JavaHost process. For example:

```
<ias-component id="coreapplication_obijh1">
```

3. Within the `environment` tag, add new variables required for Hyperion Financial Management, using the following example as a guide. Update the actual values as appropriate for your installation.

```
<variable id="HYPERION_HOME" value="C:/Hyperion"/>
<variable id="ADM_HOME" value="$HYPERION_HOME/common/ADM/9.3.1"/>
<variable id="HFM_HOME" value="$HYPERION_HOME/FinancialManagement/Client"/>
<variable id="ADM_ESS_NATIVE_HOME" value="$ADM_HOME/Essbase/9.3.1"/>
```

4. Update the `PATH` variable to include `ADM_ESS_NATIVE_HOME/bin` and `HFM_HOME`. For example:

```
<variable id="PATH" value="$ADM_ESS_NATIVE_HOME/bin:$HFM_HOME:$ORACLE_
HOME/bifoundation/server/bin:$ORACLE_HOME/bifoundation/web/bin:$ORACLE_
HOME/bifoundation/odbc/lib:$ORACLE_INSTANCE:$ORACLE_HOME/lib" append="true"/>
```

5. Save and close the file.
6. Restart OPMN.
7. Repeat these steps on each computer that runs the Oracle BI JavaHost process. If you are running multiple JavaHost instances on the same computer, be sure to update the `ias-component` tag appropriately for each instance in `opmn.xml` (for example, `ias-component id="coreapplication_obijh1"`, `ias-component id="coreapplication_obijh2"`, and so on).

You should run more than one JavaHost process to ensure that the JavaHost is not a single point of failure for HFM access. To do this, scale out the JavaHost process using Fusion Middleware Control. See "Scaling Your Deployment" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about scaling out processes.

For Hyperion Financial Management 11.1.2:

1. Open `opmn.xml` for editing. You can find `opmn.xml` at:

```
ORACLE_INSTANCE/config/OPMN/opmn/opmn.xml
```

2. Locate the `ias-component` tag for the JavaHost process. For example:

```
<ias-component id="coreapplication_obijh1">
```


3. Within the environment tag, add new variables required for Hyperion Financial Management, using the following example as a guide. Update the actual values as appropriate for your installation.

```
<variable id="HYPERION_HOME" value="C:/oracle/Middleware/EPMSysstem11R1"/>
<variable id="EPM_HOME" value="C:/MW_Home/Oracle_BI1/common/jlib/11.1.2.0"/>
<variable id="ADM_HOME" value="C:/MW_Home/Oracle_BI1/common/ADM/11.1.2.0"/>
<variable id="HFM_HOME" value="$HYPERION_
HOME/products/FinancialManagement/Client"/>
<variable id="ADM_ESS_NATIVE_HOME" value="$ADM_HOME"/>
<variable id="CSS_HOME" value="c:/MW_Home/Oracle_BI1/common/CSS/11.1.2.0"/>
```

4. Update the PATH variable to include ADM_HOME/bin and the client and common subdirectories of HYPERION_HOME. For example:

```
<variable id="PATH" value="C:\oracle\product\10.2.0\db_1\BIN;$ADM_
HOME/bin;$HYPERION_HOME/products/FinancialManagement/Common;$HYPERION_
HOME/products/FinancialManagement/Client;$ORACLE_
HOME/common/ODBC/Merant/5.3/lib;$ORACLE_HOME/bifoundation/server/bin;$ORACLE_
HOME/bifoundation/web/bin;$ORACLE_HOME/bifoundation/odbc/lib;$ORACLE_
INSTANCE;$ORACLE_HOME/lib" append="true"/>
```

5. Locate the data id="start-args" tag and update the classpath attribute to include the location of required jar files. For example:

```
<data id="start-args" value="-Xmx1024M -Djava.security.manager=default -DXDO_
FONT_DIR=$ORACLE_HOME/common/fonts -Djava.security.policy=$ORACLE_
INSTANCE/config/$COMPONENT_TYPE/$COMPONENT_NAME/java.policy
-Djava.util.logging.config.class=oracle.core.ojdl.logging.LoggingConfiguration
-Doracle.core.ojdl.logging.config.file=$ORACLE_INSTANCE/config/$COMPONENT_
TYPE/$COMPONENT_NAME/logging-config.xml -Doracle.home=$ORACLE_HOME
-Doracle.instance=$ORACLE_INSTANCE -Doracle.component.type=$COMPONENT_TYPE
-Doracle.component.name=$COMPONENT_NAME -classpath $COMMON_COMPONENTS_
HOME/modules/oracle.odl_11.1.1/ojdl.jar;$ORACLE_
HOME/bifoundation/javahost/lib/core/sautils.jar;$ORACLE_
HOME/bifoundation/javahost/lib/core/mad.jar;$ORACLE_
HOME/bifoundation/javahost/lib/core/sawconnect.jar;$ORACLE_
HOME/bifoundation/javahost/lib/core/javahost.jar;$COMMON_COMPONENTS_
HOME/modules/oracle.xdk_11.1.0/xmlparserv2.jar;$ORACLE_
HOME/opmn/lib/ons.jar;$COMMON_COMPONENTS_HOME/modules/oracle.jps_
11.1.1/jps-manifest.jar;$COMMON_COMPONENTS_HOME/modules/oracle.jps_
11.1.1/jps-api.jar;$COMMON_COMPONENTS_HOME/modules/oracle.jps_
11.1.1/jps-common.jar;$COMMON_COMPONENTS_HOME/modules/oracle.jps_
11.1.1/jps-internal.jar;$COMMON_COMPONENTS_HOME/modules/oracle.nlsrtl_
11.1.0/orai18n.jar;$COMMON_COMPONENTS_HOME/modules/oracle.nlsrtl_
11.1.0/orai18n-collation.jar;$COMMON_COMPONENTS_HOME/modules/oracle.nlsrtl_
11.1.0/orai18n-mapping.jar;$COMMON_COMPONENTS_HOME/modules/oracle.dms_
11.1.1/dms.jar;$EPM_HOME/epm.jar;$ADM_HOME/lib/adm.jar;$ADM_
HOME/lib/ap.jar;$CSS_HOME/lib/css.jar;$CSS_HOME/lib/common-model.jar$
com.siebel.analytics.javahost.standalone.Main"/>
```

6. Save and close the file.
7. Go to the directory *ORACLE_HOME/bifoundation/javahost/lib/obisintegration/adm* and delete all jar files except for *admintegration.jar* and *admimport.jar*.
8. Restart the JavaHost process.
9. Repeat these steps on each computer that runs the Oracle BI JavaHost process. If you are running multiple JavaHost instances on the same computer, be sure to

update the `ias-component` tag appropriately for each instance in `opmn.xml` (for example, `ias-component id="coreapplication_obijh1"`, `ias-component id="coreapplication_obijh2"`, and so on).

You should run more than one JavaHost process to ensure that the JavaHost is not a single point of failure for HFM access. To do this, scale out the JavaHost process using Fusion Middleware Control. See "Scaling Your Deployment" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about scaling out processes.

Setting Up SAP/BW Data Sources

You can connect to SAP/BW data sources using either the XMLA connection type, or the SAP BW Native connection type (BAPI). SAP BW Native connections might not be available on certain platforms. See "[System Requirements and Certification](#)" for more information.

To connect to SAP/BW data sources using the SAP BW Native connection type, you must first download the **BAPI Adapter for OBI Package** from the Bristlecone Web site at:

http://www.bristleconelabs.com/edel/OBIEE_BAPIAdaptor_download.asp

Then, follow the configuration instructions in the Readme provided with the download.

No preconfiguration steps are required to connect to SAP/BW over XMLA.

Setting Up Oracle RPAS Data Sources

Oracle BI Server can connect to Oracle RPAS (Retail Predictive Application Server) data sources through ODBC DSNs. To set up Oracle RPAS data sources, you must first install the Oracle RPAS ODBC driver. During set up of the ODBC DSN, you must select the **SQLExtendedFetch** option, select DBMS from the **Authentication Method** list, and select No from the **Normalize Dimension Tables** list. See "[About Importing Metadata from Oracle RPAS Data Sources](#)" for more information.

On Windows systems, you can connect to Oracle RPAS data sources for both initial import and for access to the data source during query execution. On UNIX systems, you can only connect to Oracle RPAS data sources for data access.

See "[Configuring Oracle RPAS ODBC Data Sources on AIX UNIX](#)" for information about how to enable ODBC access to Oracle RPAS data sources when the Oracle BI Server is running on AIX UNIX.

Importing Metadata from Relational Data Sources

You can import metadata for supported relational data source types by selecting one of the following import options and the appropriate connection type:

- **Local Machine.** Available in both offline and online modes. Use this option when you have all database connections set up on your computer. Note the following:
 - Most metadata imports are performed using an ODBC connection type.
 - Native database gateways for metadata import are supported for Oracle Database (using OCI), DB2 (using DB2 CLI Unicode), and XML connection types. See "[Importing Metadata from XML Data Sources](#)" for information about importing XML data using the Oracle BI Server XML gateway.

- You can use the **Find** button in the Select Metadata Objects screen of the Import Metadata Wizard to find a specific table you want to import.
- **Through the Oracle BI Server.** Available in online mode. Use this option when you want to use Oracle BI Server connections to import metadata. This option lets you use the Data Source Name (DSN) of the Oracle BI Server computer to import metadata, rather than duplicating the DSNs and connectivity information on your local computer. Available connection types are ODBC, OCI, DB2 CLI (Unicode), and XML.

You must create a DSN for the Oracle BI Server before you can import metadata through the Oracle BI Server. In addition, the Oracle BI Server must be running. For more information about the Oracle BI Server DSN, see "Integrating Other Clients with Oracle Business Intelligence" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Note: When it is running on a UNIX platform, the Oracle BI Server does not support importing metadata using an ODBC connection type.

See "[Importing Metadata from Multidimensional Data Sources](#)" and [Chapter 5, "Working with ADF Business Component Data Sources"](#) for information about importing from other data sources.

When you import physical tables, be careful to limit the import to only those tables that contain data that are likely to be used in the business models you create. You can use the Find feature to locate and select the tables that you want to import. Importing large numbers of extraneous tables and other objects adds unnecessary complexity and increases the size of the repository.

When you import metadata for most data sources, the default is to import tables, primary keys, and foreign keys. It is recommended that you import primary and foreign keys along with your tables so that the keys are automatically created in the Physical layer. If you do not import keys, you must create them manually, which can be a time-consuming process.

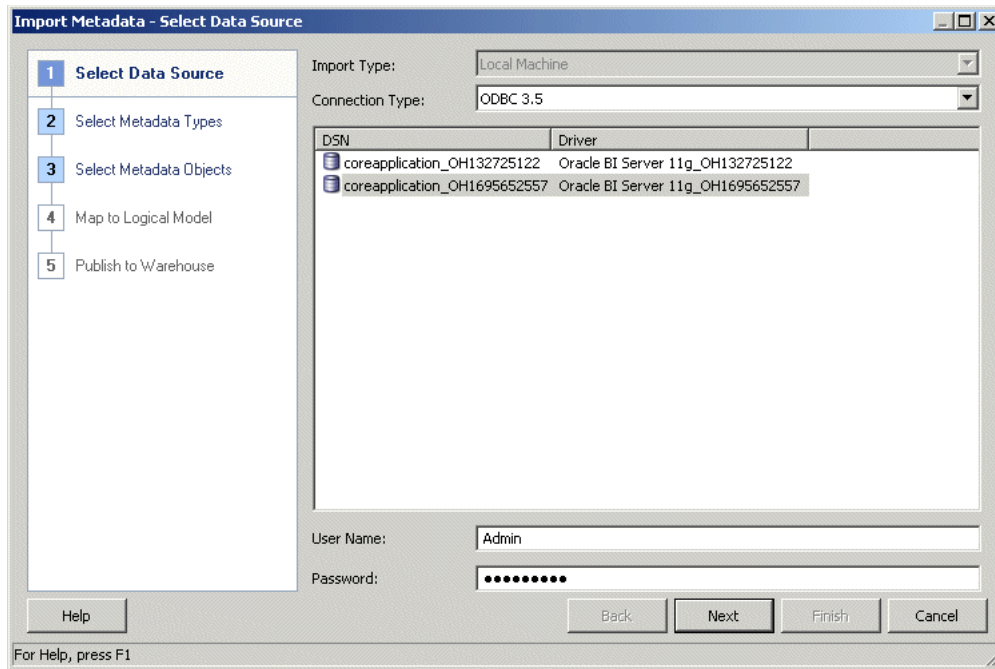
You can also import database views, aliases, synonyms, and system tables. Import these objects only if you want the Oracle BI Server to generate queries against them.

To import metadata from a relational data source:

1. In the Administration Tool, select **File**, then select **Import Metadata**. The Import Metadata Wizard appears.

Note: If you have already defined an existing database and connection pool, you can right-click the connection pool in the Physical layer and select **Import Metadata**. The Import Metadata Wizard appears with the information on the Select Data Source screen pre-filled.

[Figure 4–1](#) shows the Import Metadata Wizard.

Figure 4–1 Import Metadata Wizard: Relational Data Source

2. In the Select Data Source screen, in the **Connection Type** field, select the type of connection appropriate for your data source, such as **ODBC 3.5**.

Make sure to choose **OCI 10g/11g** if your data source is an Oracle Database. Using OCI as your connection protocol to an Oracle Database ensures better performance and provides access to native database features that are not available through ODBC.

Note: For non-Oracle databases, it is recommended that you use ODBC 3.5 or DB2 CLI (Unicode) for importing schemas with International characters, such as Japanese table and column names.

The remaining fields and options on the Select Data Source screen vary according to the connection type you selected:

- For **ODBC 2.0** and **ODBC 3.5** data sources, in the **DSN** list, select a data source from which to import the schema. Then, provide a valid user name and password for the data source.

Note that when you import through the Oracle BI Server, the DSN entries are on the Oracle BI Server, not on the local computer.

- For **OCI 10g/11g** and **DB2 CLI (Unicode)** data sources, provide the name of the data source in the **Data Source Name** field, then provide a valid user name and password for the data source.

For Oracle Database data sources, the data source name is either a full connect string or a net service name from the tnsnames.ora file. If you enter a net service name, you must ensure that you have set up a tnsnames.ora file within the Oracle Business Intelligence environment, in `ORACLE_HOME/network/admin`.

Other data source types are described in other sections:

- See ["Importing Metadata from Multidimensional Data Sources"](#) for **Essbase**, **XMLA**, **Oracle OLAP**, **Hyperion ADM**, and **SAP BW Native**. This section also describes importing from Oracle RPAS data sources over ODBC 3.5.
- See ["Importing Metadata from XML Data Sources"](#) for **XML**.
- See [Chapter 5, "Working with ADF Business Component Data Sources"](#) for **OracleADF_HTTP**.

When you have finished providing information in the Select Data Source screen, click **Next**. The Select Metadata Types screen appears.

3. Select the options for the types of objects that you want to import (for example, **Tables**, **Keys**, and **Foreign Keys**). Some options are automatically selected. Different types of data sources have different default selections, based on what is typical for that data source.

If you want to import joins, select both **Keys** and **Foreign Keys**. If you want to import system tables, you must have the system privilege for your data source. To import from Customer Relationship Management (CRM) tables, select **Metadata from CRM tables**.

4. Click **Next**. The Select Metadata Objects screen appears.
5. Select the objects you want to import in the **Available** list and move them to the **Selected** list, using the > (Import selected) or >> (Import all) buttons. You can also move objects from the **Selected** list back to the **Available** list, using the < (Remove selected) and << (Remove all) buttons.

To search for a particular item, enter a keyword in the **Find** box and then click **Find Down** or **Find Up**.

Select **Show complete structure** to view all objects, including those that have already been imported. Deselecting this option shows only the objects that are available for import. When this option is selected, objects that have already been imported appear grayed out.

6. Click **Finish**.

If some objects could not be imported, a list of warning messages appears. In the dialog displaying the messages, you can perform the following actions:

- To search for specific terms, click **Find** and then **Find Again**.
- To copy the contents of the window so that you can paste the messages in another file, click **Copy**.

After you import metadata, you should check to ensure that your database and connection pool settings are correct. In rare cases, the Oracle BI Server cannot determine the exact database type during import and instead assigns an approximate type to the database object. See ["Setting Up Database Objects"](#) and ["Creating or Changing Connection Pools"](#) for more information about working with these objects.

It is also a good practice to visually inspect the imported data in the Physical layer, such as physical columns and tables, to ensure that the import completed successfully.

About the Map to Logical Model and Publish to Warehouse Screens

The Map to Logical Model and Publish to Warehouse screens in the Import Metadata Wizard are reserved for a future release. Check the Oracle Business Intelligence chapter in *Oracle Fusion Middleware Release Notes* for updates about the functionality in these two screens.

Importing Metadata from Multidimensional Data Sources

You can use the Administration Tool to import metadata from a multidimensional data source to the Physical layer of the Oracle BI repository. The ability to use multidimensional data sources enables the Oracle BI Server to connect to sources such as Essbase, Oracle OLAP, Hyperion Financial Management, Microsoft Analysis Services, and SAP/BW (SAP/Business Warehouse) to extract data. Data from these sources can be displayed on dashboards.

The primary differences between setting up multidimensional data sources and relational data sources are in the Physical layer. The setup in the Business Model and Presentation layers for multidimensional data sources and relational data sources is almost identical.

During the import process, each cube in a multidimensional data source is created as a single physical cube table. The Oracle BI Server imports the cube metadata, including its metrics, dimensions and hierarchies. After importing the cubes, you need to ensure that the physical cube columns have the correct aggregation rule and that the hierarchy type is correct. See "[Working with Physical Hierarchy Objects](#)" for more information.

Caution: Manually creating a physical schema from a multidimensional data source is labor-intensive and error prone. Therefore, it is strongly recommended that you use the import method.

It is recommended that you remove hierarchies and columns from the Physical layer if they will not be used in the business model. This eliminates maintaining unnecessary objects in the Administration Tool and might result in better performance.

To import metadata from a multidimensional data source:

1. In the Administration Tool, select **File**, then select **Import Metadata**. The Import Metadata Wizard appears.

Note: If you have already defined an existing database and connection pool, you can right-click the connection pool in the Physical layer and select **Import Metadata**. The Import Metadata Wizard appears with the information on the Select Data Source screen pre-filled. You can also use this method to perform incremental imports.

Figure 4–2 shows the Import Metadata Wizard.

Figure 4–2 Import Metadata Wizard: Multidimensional Data Source

Note that the Map to Logical Model and Publish to Warehouse screens in the Import Metadata Wizard are reserved for a future release. Check the Oracle Business Intelligence chapter in *Oracle Fusion Middleware Release Notes* for updates about the functionality in these two screens.

2. In the Select Data Source screen, in the **Connection Type** field, select the type of connection appropriate for your data source.

The remaining fields and options on the Select Data Source screen vary according to the connection type you selected. [Table 4–1](#) describes the multidimensional connection types.

Table 4–1 Multidimensional Connection Options

Connection Type	Description
ODBC 3.5	The ODBC 3.5 connection type is used for Oracle RPAS data sources. Select the DSN entry and provide the user name and password for the selected data source. See "Setting Up ODBC Data Source Names (DSNs)" for more information.
Essbase	Provide the host name of the computer where the Essbase server is running in the Essbase Server field, then provide a valid user name and password for the data source. This information should be obtained from your data source administrator. See "Working with Essbase Data Sources" for information about how data from Essbase data sources is modeled in the Physical layer. Note: You must ensure that Essbase client libraries are installed on the computer running the Oracle BI Server before you can import metadata from Essbase data sources. See "Setting Up Essbase Data Sources" for more information.

Table 4–1 (Cont.) Multidimensional Connection Options

Connection Type	Description
XMLA	<p>The XMLA connection type is used for Microsoft Analysis Services and SAP/BW. Enter the URL of a data source from which to import the schema. Then, enter the Provider Type (such as Analysis Services 2000 or SAP/BW 3.5/7.0) and a valid user name and password for the data source.</p> <p>For Target Database, do one of the following:</p> <ul style="list-style-type: none"> ■ Select New and enter the name you want to use for the new database object. ■ Select Existing and then click Browse to select an existing database object.
Oracle OLAP	<p>Provide the name of the data source (net service name) in the Data Source Name field, then provide a valid user name and password for the data source. The data source name is the same as the entry you created in the tnsnames.ora file. Make sure that the name you use is from the tnsnames.ora file within the Oracle Business Intelligence environment. You can also choose to enter a full connect string rather than the net service name.</p> <p>For Target Database, do one of the following:</p> <ul style="list-style-type: none"> ■ Select New and enter the name you want to use for the new database object. ■ Select Existing and then click Browse to select an existing database object. <p>The JavaHost process must be running to import from Oracle OLAP data sources, for both offline and online imports.</p> <p>See "Working with Oracle OLAP Data Sources" for information about how data from Oracle OLAP data sources is modeled in the Physical layer.</p> <p>Note: Oracle Database data sources with the OLAP option can contain both relational tables and multidimensional tables. However, you should avoid having both table types in the same database object in the Administration Tool, because you may need to specify different database feature sets for the different table types.</p> <p>For example, Oracle OLAP queries fail if the database feature <code>GROUP_BY_GROUPING_SETS_SUPPORTED</code> is enabled. However, you may need this feature enabled for Oracle Database relational tables.</p> <p>As a best practice, create two separate database objects for relational and multidimensional tables.</p>

Table 4–1 (Cont.) Multidimensional Connection Options

Connection Type	Description
Hyperion ADM	<p>Provide the URL for the Hyperion Financial Management server, including the driver and application name (cube name), in the following format:</p> <pre>adm:native:HsvADMDriver:ip_or_host:application_name</pre> <p>For example:</p> <pre>adm:native:HsvADMDriver:130.35.40.80:UCFHFM</pre> <p>You also need to enter a valid user name and password for your data source.</p> <p>Note that the JavaHost process must be running to import from Hyperion Financial Management data sources, for both offline and online imports.</p> <p>See "Working with Hyperion Financial Management Data Sources" for information about how data from Hyperion Financial Management data sources is modeled in the Physical layer.</p> <p>Note: Be sure to complete the preconfiguration steps described in "Setting Up Hyperion Financial Management Data Sources" before import.</p>
SAP BW Native	<p>Provide the following information:</p> <ul style="list-style-type: none"> ■ System IP or Hostname: The host name or IP address of the SAP data server. This field corresponds to the parameter <code>ashost</code> in the SAP/BW connect string. ■ System Number: The SAP system number. This is a two-digit number assigned to an SAP instance, also called Web Application Server, or WAS. This field corresponds to the parameter <code>sysnr</code> in the SAP/BW connect string. ■ Client Number: The SAP client number. This is a three-digit number assigned to the self-contained unit called Client in SAP. A Client can be a training, development, testing, or production client, or it can represent different divisions in a large company. This field corresponds to the parameter <code>client</code> in the SAP/BW connect string. ■ Language: The SAP language code used when logging in to the data source (for example, EN for English or DE for German). This field corresponds to the parameter <code>lang</code> in the SAP/BW connect string. ■ Additional Parameters: Additional connection string parameters in the format <code>param=value</code>. Delimit multiple parameters with a colon. This field is optional. ■ User Name: A valid user name for the data source. ■ Password: The corresponding user password. The password is case-sensitive. <p>The first five fields constitute the elements of the SAP/BW connect string, in the format:</p> <pre>ashost=value:sysnr=value:client=value:lang=value:additional_param=value</pre> <p>For example:</p> <pre>ashost=10.30.0.19:sysnr=00:client=100:lang=EN</pre> <p>Note: Be sure to complete the preconfiguration steps described in "Setting Up SAP/BW Data Sources" before import.</p>

Other data source types are described in other sections:

- See ["Importing Metadata from XML Data Sources"](#) for XML.
- See ["Importing Metadata from Relational Data Sources"](#) for ODBC 2.0, OCI 10g/11g, and DB2 CLI (Unicode).
- See [Chapter 5, "Working with ADF Business Component Data Sources"](#) for OracleADF_HTTP.

When you have finished providing information in the Select Data Source screen, click **Next**.

3. For Oracle RPAS data sources only, the Select Metadata Types screen is displayed. For Oracle RPAS, select **Tables**, **Keys**, and **Foreign Keys**. Then, click **Next**.

See ["About Importing Metadata from Oracle RPAS Data Sources"](#) for more information.

4. In the Select Metadata Objects screen, select the objects you want to import in the **Available** list and move them to the **Selected** list, using the > (Import selected) or >> (Import all) buttons. You can also move objects from the **Selected** list back to the **Available** list, using the < (Remove selected) and << (Remove all) buttons.

To search for a particular item, enter a keyword in the **Find** box and then click **Find Down** or **Find Up**.

Select **Show complete structure** to view all objects, including those that have already been imported. Deselecting this option shows only the objects that are available for import. When this option is selected, objects that have already been imported appear grayed out

For Essbase data sources, select **Import UDAs** if you want to import UDAs (user-defined attributes).

5. Click **Finish**.

If some objects could not be imported, a list of warning messages appears. In the dialog displaying the messages, you can perform the following actions:

- To search for specific terms, click **Find** and then **Find Again**.
- To copy the contents of the window so that you can paste the messages in another file, click **Copy**.

After you import metadata, you should check to ensure that your database and connection pool settings are correct. In rare cases, the Oracle BI Server cannot determine the exact database type during import and instead assigns an approximate type to the database object. See ["Setting Up Database Objects"](#) and ["Creating or Changing Connection Pools"](#) for more information about working with these objects.

It is also a good practice to visually inspect the imported data in the Physical layer, such as physical columns and hierarchical levels, to ensure that the import completed successfully.

For Essbase data sources, all hierarchies are imported as Unbalanced by default. Review the **Hierarchy Type** property for each physical hierarchy and change the value if necessary. Supported hierarchy types for Essbase are **Unbalanced**, **Fully balanced**, and **Value**.

About Importing Metadata from Oracle RPAS Data Sources

This section provides important information about using the Administration Tool to import metadata from Oracle RPAS, as follows:

- Oracle RPAS schemas can only be imported on Windows.

- Before you import RPAS schemas, you must set the **Normalize Dimension Tables** field value in the ODBC DSN Setup page to **Yes** for the following reasons:
 - Setting this value to **Yes** uses an appropriate schema model (the snowflake schema) that creates joins correctly and enables drill down in the data.
 - Setting this value to **No** uses a less appropriate schema model (the star schema) that creates joins between all of the tables, causing drill down to not work correctly. Many of the joins created in this way are unwanted, and would need to be removed manually.

See ["Setting Up ODBC Data Source Names \(DSNs\)"](#) for more information.

- When you import RPAS schemas in the Administration Tool, you must import the data with joins. To do this, select the metadata types **Keys** and **Foreign Keys** in the Import Metadata Wizard.
- After you have imported RPAS schemas, you must change the **Normalize Dimension Tables** field value in the ODBC DSN Setup page back to **No**. You need to revert this setting back to **No** after import to enable the Oracle BI Server to correctly generate optimized SQL against the RPAS driver.

Note: If you do not change the **Normalize Dimension Tables** setting value to **No**, most queries will fail with an error message similar to the following:

```
[nQSError: 16001] ODBC error state: S0022 code: 0 message: [Oracle Retail][RPAS
ODBC]Column:YEAR_LABEL not found..[nQSError: 16014] SQL statement preparation
failed. Statement execute failed.
```

- If Oracle RPAS is the only data source, you must set the value of NULL_VALUES_SORT_FIRST to ON in the NQSCONFIG.INI file. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about setting values in NQSCONFIG.INI.

Importing Metadata from XML Data Sources

This section describes the use of the Extensible Markup Language (XML) as a data source. XML is the universal format for structured documents and data on the Web. It can also be used as a database to store structured data.

The Oracle BI Server supports various XML access modes, including access through the Oracle BI Server XML Gateway, as well as access through an XML ODBC driver.

This section contains the following topics:

- [About Using XML as a Data Source](#)
- [Importing Metadata from XML Data Sources Using the XML Gateway](#)
- [Importing Metadata from XML Data Sources Using XML ODBC](#)
- [Examples of XML Documents](#)

About Using XML as a Data Source

The Oracle BI Server supports the use of XML data as a data source for the Physical layer in the repository. Depending on the method used to access XML data sources, a data source may be represented by a URL pointing to one of the following sources:

- A static XML file or HTML file that contains XML data islands on the Internet (including intranet or extranet). For example:

```
tap://216.217.17.176/[DE0A48DE-1C3E-11D4-97C9-00105AA70303].XML
```

- Dynamic XML generated from a server site. For example:
tap://www.aspserver.com/example.asp
- An XML file or HTML file that contains XML data islands on a local or network drive. For example:
d:/xmldir/example.xml
d:/htmldir/island.htm

You can also specify a directory path for local or network XML files, or you can use the asterisk (*) as a wildcard with the filenames. If you specify a directory path without a filename specification (like d:/xmldir), all files with the XML suffix are imported. For example:
d:/xmldir/
d:/xmldir/exam*.xml
d:/htmldir/exam*.htm
d:/htmldir/exam*.html
- An HTML file that contains tables, defined by a pair of <table> and </table> tags. The HTML file may reside on the Internet (including intranet or extranet), or on a local or network drive. See "[Accessing HTML Tables](#)" for more information.

URLs can include repository or session variables, providing support for HTTP data sources that accept user IDs and passwords embedded in the URL. For example:

```
http://somewebserver/cgi.pl?userid=valueof(session_variable1)&password=
valueof(session_variable2)
```

This functionality also lets you create an XML data source with a location that is dynamically determined by some run-time parameters. For more information about variables, see [Chapter 18](#).

If the Oracle BI Server needs to access any nonlocal files (network files or files on the Internet, for example), you must run the Oracle BI Server using a valid user ID and password with sufficient network privileges to access these remote files.

Importing Metadata from XML Data Sources Using the XML Gateway

Using the Oracle BI Server XML Gateway, the metadata import process flattens the XML document to a tabular form using the stem of the XML filename (that is, the filename without the suffix) as the table name and the second level element in the XML document as the row delimiter. All leaf nodes are imported as columns belonging to the table. The hierarchical access path to leaf nodes is also imported.

The Oracle BI Server XML Gateway uses the metadata information contained in an XML schema. The XML schema is contained within the XML document, or is referenced within the root element of the XML document.

Where there is no schema available, all XML data is imported as text data. In building the repository, you can alter the data types of the columns in the Physical layer, overriding the data types for the corresponding columns defined in the schema. The gateway converts the incoming data to the desired type as specified in the Physical layer. You can also map the text data type to other data types in the Business Model and Mapping layer of the Administration Tool, using the CAST operator.

The Oracle BI Server XML Gateway does not support:

- Resolution of external references contained in an XML document (other than a reference to an external XML schema, as demonstrated in the example file in "Examples of XML Documents Generated by the Oracle BI Server XML Gateway").
- Element and attribute inheritance contained within the Microsoft XML schema.
- Element types of a mixed content model (such as XML elements that contain a mixture of elements and CDATA, such as `<p> hello Joe, how are you doing?</p>`).

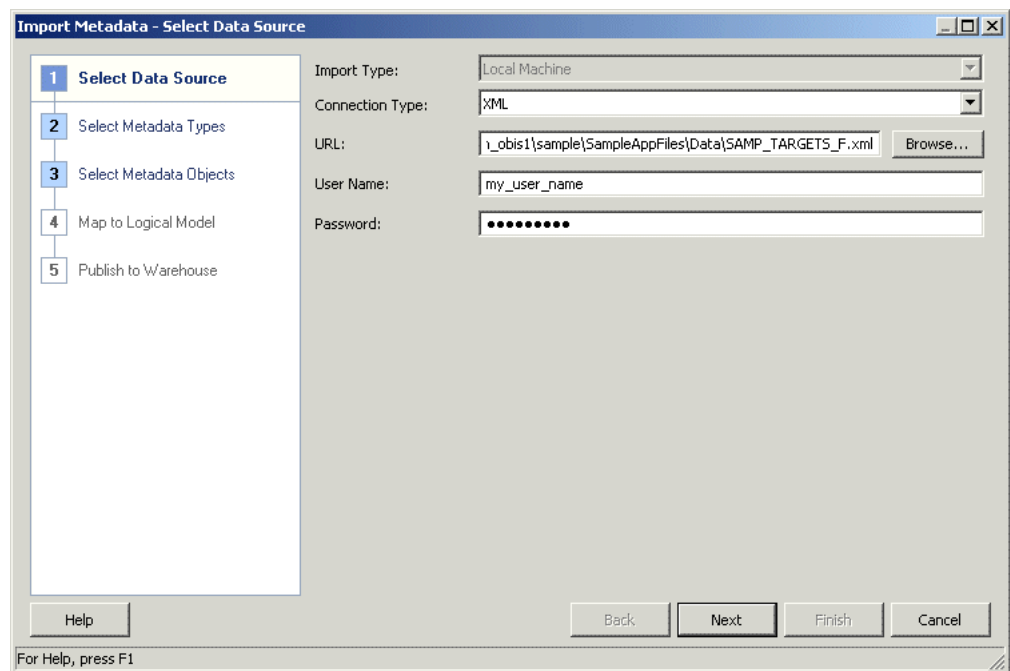
To import XML data using the Oracle BI Server XML Gateway:

1. In the Administration Tool, select **File**, then select **Import Metadata**. The Import Metadata Wizard appears.

Note: If you have already defined an existing database and connection pool, you can right-click the connection pool in the Physical layer and select **Import Metadata**. The Import Metadata Wizard appears with the information on the Select Data Source screen pre-filled.

Figure 4–3 shows the Import Metadata Wizard.

Figure 4–3 Import Metadata Wizard: XML Data Source



Note that the Map to Logical Model and Publish to Warehouse screens in the Import Metadata Wizard are reserved for a future release. Check the Oracle Business Intelligence chapter in *Oracle Fusion Middleware Release Notes* for updates about the functionality in these two screens.

2. In the Select Data Source screen, select **XML** for **Connection Type**. Then, provide the following values:
 - For **URL**, specify the XML data source URL. The Oracle BI Server XML Gateway supports all data sources described in the previous section.

URLs can include repository or session variables. For more information about variables, see [Chapter 18](#).

If you click **Browse**, the Select XML File dialog appears, from which you can select a single file. For XML documents, the file name in the URL that you specify must have the suffix `.xml`. Otherwise, the documents are treated as HTML documents.

- Type an optional user name and password in the appropriate fields for connections to HTTP sites that employ the HTTP Basic Authentication security mode.

In addition to HTTP Basic Authentication security mode, the Oracle BI Server XML Gateway also supports Secure HTTP protocol and Integrated Windows Authentication (for Windows 2000), formerly called NTLM or Windows NT Challenge/Response authentication.

When you have finished providing information in the Select Data Source screen, click **Next**. The Select Metadata Types screen appears.

3. Select the options for the types of objects that you want to import (for example, **Tables**, **Keys**, and **Foreign Keys**). The most typical options are automatically selected.

If you want to import joins, select both **Keys** and **Foreign Keys**. If you want to import system tables, you must have the system privilege for your data source.

4. Click **Next**. The Select Metadata Objects screen appears.
5. Select the objects you want to import in the **Available** list and move them to the **Selected** list, using the **>** (Import selected) or **>>** (Import all) buttons. You can also move objects from the **Selected** list back to the **Available** list, using the **<** (Remove selected) and **<<** (Remove all) buttons.

To search for a particular item, enter a keyword in the **Find** box and then click **Find Down** or **Find Up**.

Select **Show complete structure** to view all objects, including those that have already been imported. Deselecting this option shows only the objects that are available for import. When this option is selected, objects that have already been imported appear grayed out.

6. Click **Finish**.

After you import XML data, you must adjust connection pool settings. See "[Creating or Changing Connection Pools](#)" for complete information. Minimally, you can do the following:

- In the Connection Pool dialog, type a name and optional description for the connection on the General tab.
- Click the XML tab to set additional connection properties, including the URL refresh interval and the length of time to wait for a URL to load before timing out.

Because XML data sources are typically updated frequently and in real time, you can specify a refresh interval for Oracle BI Server XML Gateway data sources. The default timeout interval for queries (URL loading time-out) is 15 minutes. For more information, see "About the Refresh Interval for XML Data Sources" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Examples of XML Documents Generated by the Oracle BI Server XML Gateway

[Example 4-1](#) and [Example 4-2](#) show sample XML documents and the corresponding columns that are generated by the Oracle BI Server XML Gateway.

Example 4-1 XML Schema Contained in an External File

The following sample XML data document (mytest.xml) references an XML schema contained in an external file. The schema file is shown following the data document. The generated XML schema information available for import to the repository is shown at the end.

```
<?xml version="1.0"?>
<test xmlns="x-schema:mytest_sch.xml">

<row>
<p1>0</p1>
<p2 width="5">
  <p3>hi</p3>
  <p4>
    <p6>xx0</p6>
    <p7>yy0</p7>
  </p4>
  <p5>zz0</p5>
</p2>
</row>

<row>
<p1>1</p1>
<p2 width="6">
  <p3>how are you</p3>
  <p4>
    <p6>xx1</p6>
    <p7>yy1</p7>
  </p4>
  <p5>zz1</p5>
</p2>
</row>

<row>
<p1>a</p1>
<p2 width="7">
  <p3>hi</p3>
  <p4>
    <p6>xx2</p6>
    <p7>yy2</p7>
  </p4>
  <p5>zz2</p5>
</p2>
</row>

<row>
<p1>b</p1>
<p2 width="8">
  <p3>how are they</p3>
  <p4>
    <p6>xx3</p6>
    <p7>yy3</p7>
  </p4>
  <p5>zz2</p5>
</p2>
```

```
</row>
</test>
```

The corresponding schema file follows:

```
<Schema xmlns="urn:schemas-microsoft-com:xml-data"
  xmlns:dt="urn:schemas-microsoft-com:datatypes">
  <ElementType name="test" content="eltOnly" order="many">
    <element type="row" />
  </ElementType>
  <ElementType name="row" content="eltOnly" order="many">
    <element type="p1" />
    <element type="p2" />
  </ElementType>
  <ElementType name="p2" content="eltOnly" order="many">
    <AttributeType name="width" dt:type="int" />
    <attribute type="width" />
    <element type="p3" />
    <element type="p4" />
    <element type="p5" />
  </ElementType>
  <ElementType name="p4" content="eltOnly" order="many">
    <element type="p6" />
    <element type="p7" />
  </ElementType>
  <ElementType name="p1" content="textOnly" dt:type="string" />
  <ElementType name="p3" content="textOnly" dt:type="string" />
  <ElementType name="p5" content="textOnly" dt:type="string" />
  <ElementType name="p6" content="textOnly" dt:type="string" />
  <ElementType name="p7" content="textOnly" dt:type="string" />
</Schema>
```

The name of the table generated from the preceding XML data document (mytest.xml) would be `mytest` and the column names would be `p1`, `p3`, `p6`, `p7`, `p5`, and `width`.

In addition, to preserve the context in which each column occurs in the document and to distinguish between columns derived from XML elements with identical names but appearing in different contexts, a list of fully qualified column names is generated, based on the XPath proposal of the World Wide Web Consortium, as follows:

```
//test/row/p1
//test/row/p2/p3
//test/row/p2/p4/p6
//test/row/p2/p4/p7
//test/row/p2/p5
//test/row/p2@width
```

Example 4–2 Nested Table Structures in an XML Document

The following example is a more complex example that demonstrates the use of nested table structures in an XML document. You can optionally omit references to an external schema file, in which case all elements are treated as being of the `Varchar` character type.

```
===Invoice.xml===
<INVOICE>
  <CUSTOMER>
    <CUST_ID>1</CUST_ID>
    <FIRST_NAME>Nancy</FIRST_NAME>
    <LAST_NAME>Fuller</LAST_NAME>
    <ADDRESS>
```



```

<ADD1>507 - 20th Ave. E.,</ADD1>
<ADD2>Apt. 2A</ADD2>
<CITY>Seattle</CITY>
<STATE>WA</STATE>
<ZIP>98122</ZIP>
</ADDRESS>
<PRODUCTS>
  <CATEGORY>
    <CATEGORY_ID>CAT1</CATEGORY_ID>
    <CATEGORY_NAME>NAME1</CATEGORY_NAME>
    <ITEMS>
      <ITEM>
        <ITEM_ID>1</ITEM_ID>
        <NAME></NAME>
        <PRICE>0.50</PRICE>
        <QTY>2000</QTY>
      </ITEM>
      <ITEM>
        <ITEM_ID>2</ITEM_ID>
        <NAME>SPRITE</NAME>
        <PRICE>0.30</PRICE>
        <QTY></QTY>
      </ITEM>
    </ITEMS>
  </CATEGORY>
  <CATEGORY>
    <CATEGORY_ID>CAT2</CATEGORY_ID>
    <CATEGORY_NAME>NAME2</CATEGORY_NAME>
    <ITEMS>
      <ITEM>
        <ITEM_ID>11</ITEM_ID>
        <NAME>ACOCKE</NAME>
        <PRICE>1.50</PRICE>
        <QTY>3000</QTY>
      </ITEM>
      <ITEM>
        <ITEM_ID>12</ITEM_ID>
        <NAME>SOME SPRITE</NAME>
        <PRICE>3.30</PRICE>
        <QTY>2000</QTY>
      </ITEM>
    </ITEMS>
  </CATEGORY>
</PRODUCTS>
</CUSTOMER>
<CUSTOMER>
  <CUST_ID>2</CUST_ID>
  <FIRST_NAME>Andrew</FIRST_NAME>
  <LAST_NAME>Carnegie</LAST_NAME>
  <ADDRESS>
    <ADD1>2955 Campus Dr.</ADD1>
    <ADD2>Ste. 300</ADD2>
    <CITY>San Mateo</CITY>
    <STATE>CA</STATE>
    <ZIP>94403</ZIP>
  </ADDRESS>
  <PRODUCTS>
    <CATEGORY>
      <CATEGORY_ID>CAT22</CATEGORY_ID>
      <CATEGORY_NAME>NAMEA1</CATEGORY_NAME>

```

```

<ITEMS>
  <ITEM>
    <ITEM_ID>122</ITEM_ID>
    <NAME>DDDCOKE</NAME>
    <PRICE>11.50</PRICE>
    <QTY>2</QTY>
  </ITEM>
  <ITEM>
    <ITEM_ID>22</ITEM_ID>
    <NAME>PSPRITE</NAME>
    <PRICE>9.30</PRICE>
    <QTY>1978</QTY>
  </ITEM>
</ITEMS>
</CATEGORY>
<CATEGORY>
  <CATEGORY_ID>CAT24</CATEGORY_ID>
  <CATEGORY_NAME>NAMEA2</CATEGORY_NAME>
  <ITEMS>
    <ITEM>
      <ITEM_ID>19</ITEM_ID>
      <NAME>SOME COKE</NAME>
      <PRICE>1.58</PRICE>
      <QTY>3</QTY>
    </ITEM>
    <ITEM>
      <ITEM_ID>15</ITEM_ID>
      <NAME>DIET SPRITE</NAME>
      <PRICE>9.30</PRICE>
      <QTY>12000</QTY>
    </ITEM>
  </ITEMS>
</CATEGORY>
</PRODUCTS>
</CUSTOMER>
<CUSTOMER>
  <CUST_ID>3</CUST_ID>
  <FIRST_NAME>Margaret</FIRST_NAME>
  <LAST_NAME>Leverling</LAST_NAME>
  <ADDRESS>
    <ADD1>722 Moss Bay Blvd.</ADD1>
    <ADD2> </ADD2>
    <CITY>Kirkland</CITY>
    <STATE>WA</STATE>
    <ZIP>98033</ZIP>
  </ADDRESS>
  <PRODUCTS>
    <CATEGORY>
      <CATEGORY_ID>CAT31</CATEGORY_ID>
      <CATEGORY_NAME>NAMEA3</CATEGORY_NAME>
      <ITEMS>
        <ITEM>
          <ITEM_ID>13</ITEM_ID>
          <NAME>COKE33</NAME>
          <PRICE>30.50</PRICE>
          <QTY>20033</QTY>
        </ITEM>
        <ITEM>
          <ITEM_ID>23</ITEM_ID>
          <NAME>SPRITE33</NAME>

```

```

        <PRICE>0.38</PRICE>
        <QTY>20099</QTY>
    </ITEM>
</ITEMS>
</CATEGORY>
<CATEGORY>
    <CATEGORY_ID>CAT288</CATEGORY_ID>
    <CATEGORY_NAME>NAME H</CATEGORY_NAME>
    <ITEMS>
        <ITEM>
            <ITEM_ID>19</ITEM_ID>
            <NAME>COLA</NAME>
            <PRICE>1.0</PRICE>
            <QTY>3</QTY>
        </ITEM>
        <ITEM>
            <ITEM_ID>18</ITEM_ID>
            <NAME>MY SPRITE</NAME>
            <PRICE>8.30</PRICE>
            <QTY>123</QTY>
        </ITEM>
    </ITEMS>
</CATEGORY>
</PRODUCTS>
</CUSTOMER>
</INVOICE>

```

The generated XML schema consists of one table (INVOICE) with the following column names and their corresponding fully qualified names.

Column	Fully Qualified Name
ADD1	//INVOICE/CUSTOMER/ADDRESS/ADD1
ADD2	//INVOICE/CUSTOMER/ADDRESS/ADD2
CITY	//INVOICE/CUSTOMER/ADDRESS/CITY
STATE	//INVOICE/CUSTOMER/ADDRESS/STATE
ZIP	//INVOICE/CUSTOMER/ADDRESS/ZIP
CUST_ID	//INVOICE/CUSTOMER/CUST_ID
FIRST_NAME	//INVOICE/CUSTOMER/FIRST_NAME
LAST_NAME	//INVOICE/CUSTOMER/LAST_NAME
CATEGORY_ID	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/CATEGORY_ID
CATEGORY_NAME	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/CATEGORY_NAME
ITEM_ID	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/ITEMS/ITEM/ITEM_ID
NAME	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/ITEMS/ITEM/NAME
PRICE	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/ITEMS/ITEM/PRICE
QTY	//INVOICE/CUSTOMER/PRODUCTS/CATEGORY/ITEMS/ITEM/QTY

Only tags with values are extracted as columns. An XML query generates fully qualified tag names, to help ensure appropriate columns are retrieved.

The following shows the results of a sample query against the INVOICE table:

```
SELECT first_name, last_name, price, qty, name FROM invoice
```

```
-----
```

FIRST_NAME	LAST_NAME	PRICE	QTY	NAME
Andrew	Carnegie	1.58	3	SOME COKE
Andrew	Carnegie	11.50	2	DDDCOKE
Andrew	Carnegie	9.30	12000	DIET SPRITE
Andrew	Carnegie	9.30	1978	PSPRITE
Margar	Leverling	0.38	20099	SPRITE33
Margar	Leverling	1.0	3	COLA
Margar	Leverling	30.50	20033	COKE33
Margar	Leverling	8.30	123	MY SPRITE
Nancy	Fuller	0.30		SPRITE
Nancy	Fuller	0.50	2000	
Nancy	Fuller	1.50	3000	ACOKE
Nancy	Fuller	3.30	2000	SOME SPRITE

```
-----
```

Row count: 12

Accessing HTML Tables

The Oracle BI Server XML Gateway also supports the use of tables in HTML files as a data source. The HTML file can be identified as a URL pointing to a file on the internet (including intranet or extranet) or as a file on a local or network drive.

Even though tables, defined by the <table> and </table> tag pair, are native constructs of the HTML 4.0 specification, they are often used by Web designers as a general formatting device to achieve specific visual effects rather than as a data structure. The Oracle BI Server XML Gateway is currently the most effective in extracting tables that include specific column headers, defined by <th> and </th> tag pairs.

For tables that do not contain specific column headers, the Oracle BI Server XML Gateway employs some simple heuristics to make a best effort to determine the portions of an HTML file that appear to be genuine data tables.

The following is a sample HTML file with one table.

```
<html>
  <body>
    <table border=1 cellpadding=2 cellspacing=0>
      <tr>
        <th colspan=1>Transaction</th>
        <th colspan=2>Measurements</th>
      </tr>
      <tr>
        <th>Quality</th>
        <th>Count</th>
        <th>Percent</th>
      </tr>
      <tr>
        <td>Failed</td>
        <td>66,672</td>
        <td>4.1%</td>
      </tr>
      <tr>
        <td>Poor</td>
```

```

        <td>126,304</td>
        <td>7.7%</td>
    </tr>
    <tr>
        <td>Warning</td>
        <td>355,728</td>
        <td>21.6%</td>
    </tr>
    <tr>
        <td>OK</td>
        <td>1,095,056</td>
        <td>66.6%</td>
    </tr>
    <tr>
        <td colspan=1>Grand Total</td>
        <td>1,643,760</td>
        <td>100.0%</td>
    </tr>
</table>
</body>
</html>

```

The table name is derived from the HTML filename, and the column names are formed by concatenating the headings (defined by the <th> and </th> tag pairs) for the corresponding columns, separated by an underscore.

Assuming that our sample file is named 18.htm, the table name would be 18_0 (because it is the first table in that HTML file), with the following column names and their corresponding fully qualified names:

Column	Fully Qualified Name
Transaction_Quality	\\18_0\Transaction_Quality
Measurements_Count	\\18_0\Measurements_Count
Measurements_Percent	\\18_0\Measurements_Percent

If the table column headings appear in more than one row, the column names are formed by concatenating the corresponding field contents of those header rows.

For tables without any heading tag pairs, the Oracle BI Server XML Gateway assumes the field values (as delimited by the <td> and </td> tag pairs) in the first row to be the column names. The columns are named by the order in which they appear (c0, c1, and so on).

See "[Importing Metadata from XML Data Sources Using XML ODBC](#)" and "[Examples of XML Documents](#)" for additional XML examples.

Importing Metadata from XML Data Sources Using XML ODBC

Using the XML ODBC database type, you can access XML data sources through an ODBC interface. The data types of the XML elements representing physical columns in physical tables are derived from the data types of the XML elements as defined in the XML schema. In the absence of a proper XML schema, the default data type of string is used. Data Type settings in the Physical layer do not override those defined in the XML data sources. When accessing XML data without XML schema, use the `CAST` operator to perform data type conversions in the Business Model and Mapping layer of the Administration Tool.

To import XML data using ODBC:

1. To access XML data sources through ODBC, you first need to license and install an XML ODBC driver.
2. Create ODBC DSNs that point to the XML data sources you want to access, making sure you select the XML ODBC database type.
3. In the Administration Tool, select **File**, then select **Import Metadata**.
4. Follow the instructions in the dialogs to import the ODBC DSNs into the repository. See ["Importing Metadata from Relational Data Sources"](#) for more information.

Caution: Due to XML ODBC limitations, you must select the **Synonyms** option in the Select Metadata Types screen, or no tables are imported.

Example of an XML ODBC Data Source

[Example 4-3](#) shows an example of an XML ODBC data source in the Microsoft ADO persisted file format. Both the data and the schema could be contained inside the same document.

Example 4-3 XML ODBC Example

```
<xml xmlns:s='uuid:BDC6E3F0-6DA3-11d1-A2A3-00AA00C14882'
      xmlns:dt='uuid:C2F41010-65B3-11d1-A29F-00AA00C14882'
      xmlns:rs='urn:schemas-microsoft-com:rowset'
      xmlns:z='#RowsetSchema'>
<s:Schema id='RowsetSchema'>
  <s:ElementType name='row' content='eltOnly' rs:CommandTimeout='30'
rs:updatable='true'>
    <s:AttributeType name='ShipperID' rs:number='1' rs:writeunknown='true'
rs:basecatalog='Paint' rs:basetable='Shippers'
      rs:basecolumn='ShipperID'>
      <s:datatype dt:type='i2' dt:maxLength='2' rs:precision='5'
rs:fixedlength='true' rs:benull='false' />
    </s:AttributeType>
    <s:AttributeType name='CompanyName' rs:number='2' rs:writeunknown='true'
rs:basecatalog='Paint' rs:basetable='Shippers'
      rs:basecolumn='CompanyName'>
      <s:datatype dt:type='string' rs:dbtype='str' dt:maxLength='40'
rs:benull='false' />
    </s:AttributeType>
    <s:AttributeType name='Phone' rs:number='3' rs:nullable='true'
rs:writeunknown='true' rs:basecatalog='Paint'
      rs:basetable='Shippers' rs:basecolumn='Phone'>
      <s:datatype dt:type='string' rs:dbtype='str' dt:maxLength='24'
rs:fixedlength='true' />
    </s:AttributeType>
    <s:extends type='rs:rowbase' />
  </s:ElementType>
</s:Schema>
<rs:data>
  <z:row ShipperID='1' CompanyName='Speedy Express' Phone='(503)
555-9831' />
  <z:row ShipperID='2' CompanyName='United Package' Phone='(503)
555-3199' />
  <z:row ShipperID='3' CompanyName='Federal Shipping' Phone='(503)
```

```
555-9931          '/>
</rs:data>
</xml>
```

Examples of XML Documents

The following XML documents provide examples of several different situations and explain how the Oracle BI Server XML access method handles those situations.

- The XML documents **83.xml** and **8_sch.xml** (shown in [Example 4-4](#) and [Example 4-5](#)) demonstrate the use of the same element declarations in different scope. For example, <p3> could appear within <p2> as well as within <p4>.

Because the element <p3> in the preceding examples appears in two different scopes, each element is given a distinct column name by appending an index number to the second occurrence of the element during the import process. In this case, the second occurrence becomes p3_1. If <p3> occurs in additional contexts, they become p3_2, p3_3.

- The XML documents **83.xml** and **84.xml** (shown in [Example 4-4](#) and [Example 4-6](#)) demonstrate that multiple XML files can share the same schema (8_sch.xml).
- Internet Explorer version 5 and higher supports HTML documents containing embedded XML fragments called XML islands. The XML document **island2.htm** (shown in [Example 4-7](#)) demonstrates a simple situation where multiple XML data islands, and therefore multiple tables, could be generated from one document. One table is generated for each instance of an XML island. Tables are distinguished by appending an appropriate index to the document name. For island2.htm, the two XML tables generated would be island2_0 and island2_1.

Example 4-4 83.xml

```
===83.xml===
<?xml version="1.0"?>
<test xmlns="x-schema:8_sch.xml">|
<row>
<p1>0</p1>
<p2 width="5" height="2">
  <p3>hi</p3>
  <p4>
    <p3>hi</p3>
    <p6>xx0</p6>
    <p7>yy0</p7>
  </p4>
  <p5>zz0</p5>
</p2>
</row>

<row>
<p1>1</p1>
<p2 width="6" height="3">
  <p3>how are you</p3>
  <p4>
    <p3>hi</p3>
    <p6>xx1</p6>
    <p7>yy1</p7>
  </p4>
  <p5>zz1</p5>
</p2>
</row>
```

```
</test>
```

Example 4-5 8_sch.xml

```
===8_sch.xml===
<Schema xmlns="urn:schemas-microsoft-com:xml-data"
xmlns:dt="urn:schemas-microsoft-com:datatypes">
  <AttributeType name="height" dt:type="int" />
  <ElementType name="test" content="eltOnly" order="many">
    <AttributeType name="height" dt:type="int" />
    <element type="row"/>
  </ElementType>
  <ElementType name="row" content="eltOnly" order="many">
    <element type="p1"/>
    <element type="p2"/>
  </ElementType>
  <ElementType name="p2" content="eltOnly" order="many">
    <AttributeType name="width" dt:type="int" />
    <AttributeType name="height" dt:type="int" />
    <attribute type="width" />
    <attribute type="height" />
    <element type="p3"/>
    <element type="p4"/>
    <element type="p5"/>
  </ElementType>
  <ElementType name="p4" content="eltOnly" order="many">
    <element type="p3"/>
    <element type="p6"/>
    <element type="p7"/>
  </ElementType>
  <ElementType name="test0" content="eltOnly" order="many">
    <element type="row"/>
  </ElementType>
  <ElementType name="p1" content="textOnly" dt:type="string" />
  <ElementType name="p3" content="textOnly" dt:type="string" />
  <ElementType name="p5" content="textOnly" dt:type="string" />
  <ElementType name="p6" content="textOnly" dt:type="string" />
  <ElementType name="p7" content="textOnly" dt:type="string" />
</Schema>
```

Example 4-6 84.xml

```
===84.xml===
<?xml version="1.0"?>
<test0 xmlns="x-schema:8_sch.xml">
  <row>
    <p1>0</p1>
    <p2 width="5" height="2">
      <p3>hi</p3>
      <p4>
        <p3>hi</p3>
        <p6>xx0</p6>
        <p7>yy0</p7>
      </p4>
      <p5>zz0</p5>
    </p2>
  </row>

  <row>
    <p1>1</p1>
    <p2 width="6" height="3">
```



```

    <p3>how are you</p3>
    <p4>
      <p3>hi</p3>
      <p6>xx1</p6>
      <p7>yy1</p7>
    </p4>
    <p5>zz1</p5>
  </p2>
</row>
</test0>

```

Example 4-7 Island2.htm

```

===island2.htm===
<HTML>
  <HEAD>
<TITLE>HTML Document with Data Island</TITLE>
</HEAD>
  <BODY>
<p>This is an example of an XML data island in I.E. 5</p>
  <XML ID="12345">
    <test>
      <row>
        <field1>00</field1>
        <field2>01</field2>
      </row>
      <row>
        <field1>10</field1>
        <field2>11</field2>
      </row>
      <row>
        <field1>20</field1>
        <field2>21</field2>
      </row>
    </test>
  </XML>
<p>End of first example.</p>
  <XML ID="12346">
    <test>
      <row>
        <field11>00</field11>
        <field12>01</field12>
      </row>
      <row>
        <field11>10</field11>
        <field12>11</field12>
      </row>
      <row>
        <field11>20</field11>
        <field12>21</field12>
      </row>
    </test>
  </XML>
<p>End of second example.</p>
</BODY>
</HTML>

```

Using a Standby Database with Oracle Business Intelligence

A standby database is used mainly for its high availability and failover functions as a backup for the primary database. In a standby database configuration, there is regularly scheduled replication from the primary database to the secondary database. The latency of this replication must be short enough that writing to the primary database while reading from the secondary database does not cause any synchronization or data integrity problems.

Because a standby database is essentially a read-only database, it can be used as a business intelligence query server, relieving the workload of the primary database and improving query performance.

The following sections explain how to use a standby database with Oracle Business Intelligence:

- [About Using a Standby Database with Oracle Business Intelligence](#)
- [Creating the Database Object for the Standby Database Configuration](#)
- [Creating Connection Pools for the Standby Database Configuration](#)
- [Updating Write-Back Scripts in a Standby Database Configuration](#)
- [Setting Up Usage Tracking in a Standby Database Configuration](#)
- [Setting Up Event Polling in a Standby Database Configuration](#)
- [Setting Up Oracle BI Scheduler in a Standby Database Configuration](#)

About Using a Standby Database with Oracle Business Intelligence

In a standby database configuration, you have two databases: a primary database that handles all write operations and is the source of truth for data integrity, and a secondary database that is exposed as a read-only source. When you use a standby database configuration with Oracle Business Intelligence, all write operations are offloaded to the primary database, and read operations are sent to the standby database.

Write operations that need to be routed to the primary source may include the following:

- Oracle BI Scheduler job and instance data
- Temporary tables for performance enhancements
- Writeback scripts for aggregate persistence
- Usage tracking data, if usage tracking has been enabled
- Event polling table data, if event polling tables are being used

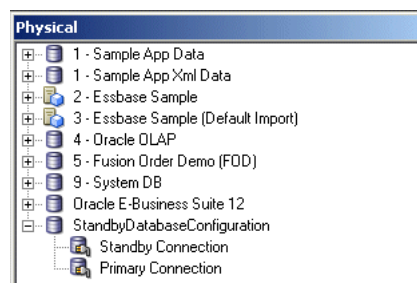
The following list provides an overview of how to configure the Oracle BI Server to use a standby database:

1. Create a single database object for the standby database configuration, with temporary table creation disabled.
2. Configure two connection pools for the database object:
 - A read-only connection pool that points to the standby database
 - A second connection pool that points to the primary database for write operations

3. Update any connection scripts that write to the database so that they explicitly specify the primary database connection pool.
4. If usage tracking has been enabled, update the usage tracking configuration to use the primary connection.
5. If event polling tables are being used, update the event polling database configuration to use the primary connection.
6. Ensure that Oracle BI Scheduler is not configured to use any standby sources.

Even though there are two separate physical data sources for the standby database configuration, you create only one database object in the Physical layer. [Figure 4–4](#) shows the database object and connection pools for the standby database configuration in the Physical layer.

Figure 4–4 Standby Database Configuration in the Physical Layer



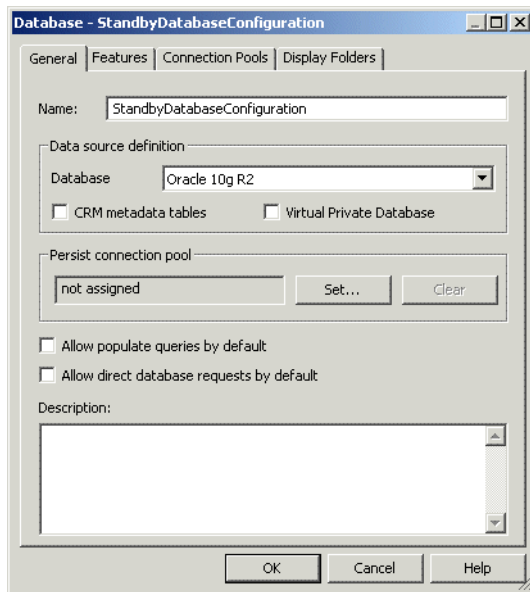
Creating the Database Object for the Standby Database Configuration

Use the Administration Tool to create a database object in the repository for the standby database configuration. When you create the database object, make sure that the persist connection pool is not assigned, to prevent the Oracle BI Server from creating temporary tables in the standby database.

To create a database object for the standby database configuration:

1. In the Administration Tool, right-click the Physical layer and select **New Database**.
2. Provide a name for the database, and then select the type of database in the Database list.
3. Ensure that the **Persist connection pool** is "not assigned."

[Figure 4–5](#) shows the Database dialog for a standby database configuration.

Figure 4–5 Database Dialog for Standby Database Configuration

Creating Connection Pools for the Standby Database Configuration

After you have created a database object in the repository for the standby database configuration, use the Administration Tool to create two connection pools: one that points to the standby database, and another that points to the primary database.

Because the standby connection pool is used for the majority of connections, make sure that the standby connection pool is listed first.

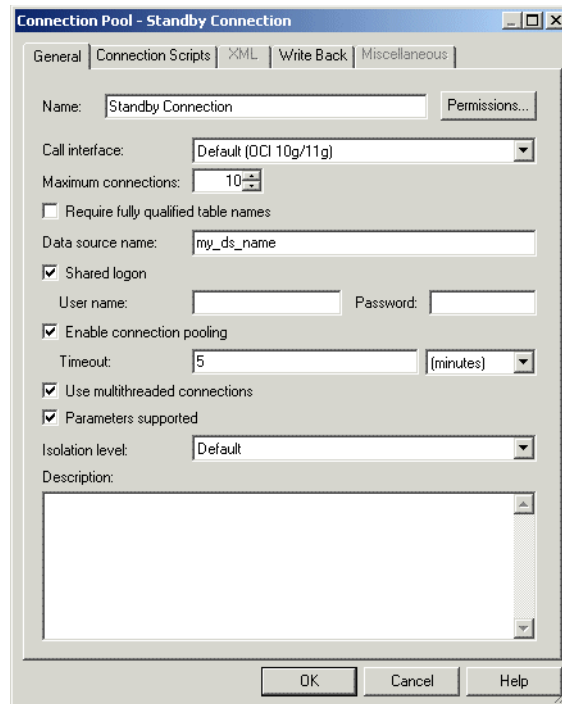
Note: Connection pools are used in the order listed, until the maximum number of connections is achieved. Ensure that the maximum number of connections is set in accordance with the standby database tuning.

See "[Creating or Changing Connection Pools](#)" for more information about setting the maximum number of connections.

To create a standby connection pool for the standby database configuration:

1. In the Administration Tool, in the Physical layer, right-click the database object for the standby database configuration and select **New Object**, then select **Connection Pool**.
2. Provide a name for the connection pool, and ensure that the call interface is appropriate for the standby database type.
3. Provide the **Data source name** for the standby database.
4. Enter a user name and password for the standby database.
5. Click **OK**.

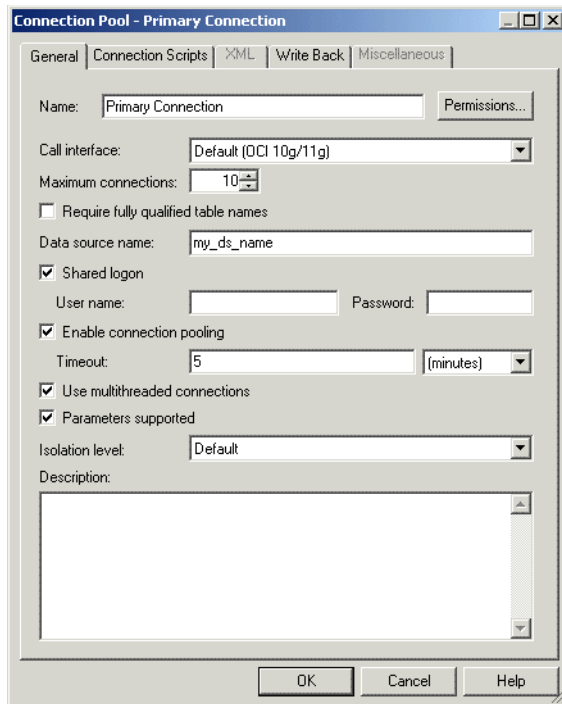
[Figure 4–6](#) shows the Connection Pool dialog for the standby connection pool.

Figure 4–6 Connection Pool Dialog for Standby Connection Pool

To create the primary connection pool for the standby database configuration:

1. In the Administration Tool, in the Physical layer, right-click the database object for the standby database configuration and select **New Object**, then select **Connection Pool**.
2. Provide a name for the connection pool, and ensure that the call interface is appropriate for the primary database type.
3. Provide the **Data source name** for the primary database.
4. Enter a user name and password for the primary database.
5. Click **OK**.

Figure 4–7 shows the Connection Pool dialog for the primary connection pool.

Figure 4–7 Connection Pool Dialog for Primary Connection Pool

Updating Write-Back Scripts in a Standby Database Configuration

If you use scripts that write to the database, such as scripts for aggregate persistence, you must update the scripts to explicitly refer to the primary connection pool. Information written through the primary connection will be automatically transferred to the standby database (through the regularly scheduled replication between the primary and secondary databases), and will become available through the standby connection pool.

The following example shows a writeback script for aggregate persistence that explicitly specifies the primary connection pool:

```
create aggregates sc_rev_qty_yr_cat for "DimSnowflakeSales"."SalesFacts"
("Revenue", "QtySold") at levels ("DimSnowflakeSales"."Time"."Year",
"DimSnowflakeSales"."Product"."Category") using connection pool
"StandbyDemo"."Primary Connection" in "StandbyDemo"."My_Schema"
```

Setting Up Usage Tracking in a Standby Database Configuration

The Oracle BI Server supports the collection of usage tracking data. When usage tracking is enabled, the Oracle BI Server collects usage tracking data for each query and writes statistics to a usage tracking log file or inserts them directly to a database table.

If you want to enable usage tracking on a standby database configuration using direct insertion, you must create the table used to store the usage tracking data (typically S_NQ_ACCT) on the primary database. Then, import the table into the physical layer of the repository using the Administration Tool.

You must ensure that the database object for the usage tracking table is configured with both the standby connection pool and the primary connection pool. Then, ensure that the CONNECTION_POOL parameter for usage tracking points to the primary database. For example, in NQSConfig.ini:

```
CONNECTION_POOL = "StandbyDatabaseConfiguration"."Primary  
Connection";
```

See "Managing Usage Tracking" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for full information about usage tracking, including how to create tables for usage tracking data and how to set parameters for usage tracking.

Setting Up Event Polling in a Standby Database Configuration

You can use an Oracle BI Server event polling table (event table) as a way to notify the Oracle BI Server that one or more physical tables have been updated. The event table is a physical table that resides on a database accessible to the Oracle BI Server. It is normally exposed only in the Physical layer of the Administration Tool, where it is identified in the Physical Table dialog as an Oracle BI Server event table.

The Oracle BI Server requires write access to the event polling table. Because of this, if you are using event polling in a standby database configuration, you must ensure that the database object for the event table only references the primary connection pool.

See "Cache Event Processing with an Event Polling Table" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for full information about event polling, including how to set up, activate, and populate event tables.

Setting Up Oracle BI Scheduler in a Standby Database Configuration

Oracle BI Scheduler is an extensible application and server that manages and schedules jobs, both scripted and unscripted. To use Oracle BI Scheduler in a standby database configuration, you must ensure that the database object for Oracle BI Scheduler only references the primary connection pool.

See "Configuration Tasks for Oracle BI Scheduler" in *Oracle Fusion Middleware Scheduling Jobs Guide for Oracle Business Intelligence Enterprise Edition* for full information about setting up and using Oracle BI Scheduler.

Working with ADF Business Component Data Sources

Oracle Business Intelligence supports connecting to ADF Business Components as data sources. This enables Oracle Business Intelligence users to query data from any application that is built using the ADF Framework. For example, because Oracle CRM applications are developed using the ADF Framework, Oracle Business Intelligence users can report on CRM data using an ADF Business Component data source that implements the required ADF Application Programming Interface (API).

By using the ADF components as a data source to the Oracle BI Server, users can quickly integrate operational reporting with any application that is built on top of the ADF Framework.

This chapter contains the following topics:

- [What Are ADF Business Components?](#)
- [What Happens During Import?](#)
- [About Specifying a SQL Bypass Database](#)
- [Setting Up ADF Business Component Data Sources](#)
- [Importing Metadata from ADF Business Component Data Sources](#)
- [Enabling the Ability to Pass Custom Parameters to the ADF Application](#)
- [Propagating Labels and Tooltips from ADF Business Component Data Sources](#)

What Are ADF Business Components?

Oracle Application Development Framework (Oracle ADF) is an object-relational framework that can be used to create J2EE business services and expose underlying database objects. This framework provides an abstraction layer that enables application developers to build applications quickly and efficiently.

When you use Oracle ADF to build service-oriented Java EE applications, you implement your core business logic as one or more business services. These back-end services provide clients with a way to query, insert, update, and delete business data as required, while enforcing appropriate business rules. ADF Business Components are prebuilt application objects that provide a ready-to-use implementation of Java EE design patterns and best practices.

The ADF model is represented through the ADF Business Component constructs called Entity Objects and View Objects, usually constructed and defined during design time:

- **Entity Objects:** ADF framework components that represent a row in a database table and simplify modifying its data. Importantly, it enables you to encapsulate domain business logic for those rows to ensure your business policies and rules are consistently validated.
- **View Objects:** ADF framework components that encapsulate a SQL query and simplify working with its results. In addition to *read-only view objects*, there are *entity-based view objects* that support updatable rows. The view object queries just the data needed for the client-facing task at hand, then cooperates with one or more entity objects in your business domain layer to automatically validate and save changes made to its view rows. Like the read-only view object, an entity-based view object encapsulates a SQL query, can be linked into master/detail hierarchies using view links, and can be used in the data model of your application modules.

Applications built using ADF obtain their data by querying the defined View Objects using the ADF APIs.

The ADF model also includes an *application module*, which is the transactional component that UI clients use to work with application data. It defines an updatable data model along with top-level procedures and functions (called service methods) related to a logical unit of work related to an end-user task.

The application module serves as a container for multiple View Objects and Entity Objects, and also contains configuration related to the JDBC data source.

About Operational Reporting with ADF Business Components

You can use Oracle Business Intelligence integration with ADF Business Components to generate reports on data within your applications. For example, you can generate reports based on expense report data entered into an Expense Application.

To do this, you would first import the Expense Application metadata into the Oracle BI repository using the Administration Tool, then map the data from the Physical layer to the Business Model and Mapping layer and Presentation layer. After you restart the Oracle BI Server and reload the metadata into Oracle BI Presentation Services, you can log in to Oracle BI Answers and drag and drop the columns to generate a report on the Expense Application data. For example, you can select columns to view a report of your expenses grouped by category.

What Happens During Import?

On import, the required physical tables and complex joins are automatically created. The instances (ViewObject and ViewLink) are imported into Oracle Business Intelligence. During query execution, the definitions retrieved from these instances are used to create the CompositeVO in ADF.

These complex joins are 'dummy joins' and are not executed in Oracle Business Intelligence. Instead, they denote ViewLink instances that connect pairs of View Objects in the ADF model. The physical table and complex join names correspond to the fully qualified ViewObject and ViewLink instance names, respectively. This convention allows arbitrary nesting of ApplicationModules in the ADF model.

Note that the **External Expression** field in the Complex Join dialog for ADF data sources shows an arbitrary expression that has no meaning. This field is reserved for a future release.

The name of the automatically generated joins follow a naming convention similar to *ViewObjectName1_ViewObjectName2* (for example, *AppModuleAM.AP_VO1_*

AppModuleAM_BU_VO1). The ViewLink instance name appears in the **ViewLink Name** field of the Complex Join dialog.

The complex joins are only created automatically if a ViewLink instance is available. They are not created for ViewLink definitions. Joins using ViewLink definitions must be created manually. To do this, specify the ViewLink definition name in the **ViewLink Name** field of the Complex Join dialog.

Alternatively, Oracle Business Intelligence joins between VOs in different ApplicationModules are created upon import from ADF if custom properties are defined on the ApplicationModule. Note the following:

- The property name format is *BI_VIEW_LINK_property_name*
- The property value format is *source_VO_instance_name, ViewLink_definition_name, destination_VO_instance_name*

Be sure to use the fully qualified VO instance names for the source and destination VOs, as well as the fully qualified package name for the ViewLink definition.

About Specifying a SQL Bypass Database

The Oracle BI Server can automatically create composite View Objects at run time, so that an ad-hoc BI query can reference multiple View Objects in the ADF layer. For improved performance, a SQL bypass query is generated that incorporates the projection columns, filters, and joins required by the BI query.

The SQL Bypass feature directly queries the database so that aggregations and other transformations are pushed down where possible, reducing the amount of data streamed and worked on in Oracle Business Intelligence. When using a SQL Bypass database, the Oracle BI Server gets the VO query from the ADF Business Component data source and then wraps it with the aggregations in the Logical SQL query. The query, including the aggregations, is then executed in the database. Because the database computes the aggregation and less rows are streamed back to Oracle Business Intelligence, using a SQL Bypass database can result in significant performance gains.

Multiple View Objects are modeled as separate BI physical tables and are connected with dummy complex joins. These joins only represent the ViewLinks in the ADF model and are not executed by the Oracle BI Server.

You can specify the name of the SQL Bypass database in the connection pool for the ADF Business Component data source. The SQL Bypass database must be a physical database in the Physical layer of the repository. The database object for the SQL Bypass database must have a valid connection pool, with connection information that points to the same database that is being used by the JDBC Data source defined in the Oracle WebLogic Server that runs the ADF application.

The SQL Bypass database does not need to have any tables under it. After a valid database name is supplied, the SQL Bypass feature is enabled for all queries against that ADF database.

Setting Up ADF Business Component Data Sources

This section explains how to configure your ADF Business Components for use with Oracle Business Intelligence.

See "[System Requirements and Certification](#)" for information about supported versions.

This section contains the following topics:

- [Creating a WebLogic Domain](#)
- [Deploying OBIEEBroker as a Shared Library in Oracle WebLogic Server](#)
- [Deploying the Application EAR File to Oracle WebLogic Server from JDeveloper](#)
- [Setting Up a JDBC Data Source in the WebLogic Server](#)
- [Setting the Logging Level for the Deployed Application in Oracle WebLogic Server](#)

Creating a WebLogic Domain

Create a WebLogic Domain for your ADF Business Components that supports WebLogic Server, Oracle Application Core (Webapp), and Oracle JRF.

To create a WebLogic domain that supports the required components:

1. Start the WebLogic Configuration Wizard. For example, on Windows, run `MW_HOME\wlserver_10.3\common\bin\config.cmd`.
2. Select **Create a new WebLogic domain** and click **Next**.
3. On the Select Domain Source screen, ensure that **Basic WebLogic Server Domain**, **Oracle JRF**, and **Oracle Application Core (Webapp)** are selected.
4. Follow the remaining steps in the wizard, providing values appropriate for your environment.
5. Click **Create** on the Configuration Summary screen to create the domain.

You can start and stop the Oracle WebLogic Server for this domain using command-line scripts in the domain directory. For example, on Windows, use the following:

- `MW_HOME\user_projects\domains\domain_name\bin\startWebLogic.cmd`
- `MW_HOME\user_projects\domains\domain_name\bin\stopWebLogic.cmd`

Deploying OBIEEBroker as a Shared Library in Oracle WebLogic Server

The OBIEEBroker shared library is installed as part of your Oracle Business Intelligence installation. You need to deploy the OBIEEBroker library as a shared library in Oracle WebLogic Server by installing it (making its physical file or directory known to Oracle WebLogic Server) and starting it. After the library has been installed and started, other deployed modules can reference the library.

To deploy OBIEEBroker as a shared library in Oracle WebLogic Server:

1. Ensure that Oracle WebLogic Server is running. If it is not running, start it. For example, on Windows, run `MW_HOME\user_projects\domains\your_domain\bin\startWebLogic.cmd`.
2. Open the WebLogic Server Administration Console. For example, if your Oracle WebLogic Server is running locally on port 7001, go to `http://localhost:7001/console`.
3. Log in to the WebLogic Server Administration Console with the credentials you created when you set up your WebLogic domain.
4. In the Change Center, click **Lock & Edit**.
5. On the Home Page, in the left pane, click **Deployments**.
6. In the right pane, click **Install**.

7. Using the Install Application Assistant, locate the OBIEEBroker EAR file. You can find this file at:

```
ORACLE_HOME\bifoundation\jvahost\lib\obisintegration\adf\
oracle.bi.integration.adf.ear
```

8. Click **Next**.
9. Select **Install this deployment as a library** and click **Next**.
10. Select the servers and/or clusters to which you want to deploy the OBIEEBroker library. Make sure to select all servers and clusters to which modules or applications that reference the library are deployed.
11. Click **Next**.
12. You can optionally update settings about the deployment. Typically, the default values are adequate. Click **Help** for more information.
13. Click **Next**, then click **Finish** to complete the installation.
14. In the Change Center, click **Activate Changes**.

Deploying the Application EAR File to Oracle WebLogic Server from JDeveloper

Follow the steps in this section to deploy the application EAR file to Oracle WebLogic Server from JDeveloper. Before beginning this procedure, ensure that the following conditions are true:

- You have an ADF Model project that contains AMs and VOs that will be exposed to Oracle Business Intelligence.
- You have deployed OBIEEBroker as a shared library in Oracle WebLogic Server. See "[Deploying OBIEEBroker as a Shared Library in Oracle WebLogic Server](#)" for more information.
- Oracle WebLogic Server is running.

To deploy the application EAR file to Oracle WebLogic Server from JDeveloper:

1. Start JDeveloper. For example, on Windows, run `MW_HOME\jdeveloper\jdev\bin\jdev.exe`.
2. Select **File**, then select **Open** to open the project that contains your ADF Business Components in JDeveloper. If prompted, allow JDeveloper to migrate the project to the latest version.
3. Create a new Application Module configuration, as follows:
 - a. In the Model project, double click the application module, then click the Configurations tab for that application module.
 - b. Create a new configuration with the following characteristics:
 - Select **JDBC DataSource** for **Connection Type**.
 - Keep the default **DataSource Name** (for example, `java:comp/env/jdbc/ApplicationDBDS`).

When you set up the JDBC data source in Oracle WebLogic Server in a later step, you use part of this DataSource Name as the JNDI name required by Oracle WebLogic Server. The JNDI name is the DataSource Name without the `java:comp/env` context prefix (for example, `jdbc/ApplicationDBDS`).

4. Create a Business Component Archive deployment provide, as follows:

- a. In the Projects window, right-click the Model project and choose **New**.
 - b. Select **Deployment Profiles** under **General** in the left pane, then choose **Business Components Archive** in the right pane and click **OK**.
 - c. Provide a name for the deployment profile (for example, MyApplication_Archive) and click **OK**.
 - d. On the Deployment page, click **OK**.
5. In the Projects window, right-click the Model project and select **Deploy > your_deployment_profile_name > Deploy**, or use the deployment wizard by selecting **Deploy to File**.

After the project has been deployed, two jar files are created in the deploy directory for the Model project (for example, MyApplication_Archive_Common.jar and MyApplication_Archive_MiddleTier.jar).

6. Create a new Web Project for the application, as follows:
 - a. Right-click the global application and select **New Project**.
 - b. Select **Projects** from the left pane, then select **Web Project** from the right pane.
 - c. Provide a project name (for example, **OBIEEBroker**).
 - d. Click **Next** until you reach the **Web Project Profile** page.
 - e. Modify the **Java EE Context Root** to a name that better represents your application (for example, MyApplication).

This value determines the URL that you use to connect to the application from Oracle Business Intelligence (for example, `http://localhost:7001/MyApplication/obieebroker`).
7. Edit the Profile Dependencies of the WAR deployment, as follows:
 - a. Right-click the Web Project you just created (for example, **OBIEEBroker**) and select **Project Properties**.
 - b. From the left pane, select **Deployment**. Then, open the WAR File deployment profile on the right pane.
 - c. Select **Profile Dependencies** from the left pane. Then, on the right pane, select the Common and MiddleTier deployment profiles of your Model project.

Following this step ensures that the Business Component Archives for the Model project are included in the WAR file.
8. Expand the Web Project and open web.xml. Then, go to the source view of the file.
9. In the web.xml source, replace the content within the <web-app> element with the following:

```
<filter>
  <filter-name>ServletADFFilter</filter-name>
  <filter-class>oracle.adf.share.http.ServletADFFilter</filter-class>
</filter>
<filter-mapping>
  <filter-name>ServletADFFilter</filter-name>
  <servlet-name>OBIEEBroker</servlet-name>
  <dispatcher>FORWARD</dispatcher>
  <dispatcher>REQUEST</dispatcher>
</filter-mapping>
<servlet>
  <servlet-name>OBIEEBroker</servlet-name>
```

```

    <servlet-class>oracle.bi.integration.adf.v11g.obieebroker.OBIEEBroker
  </servlet-class>
</servlet>
<servlet-mapping>
  <servlet-name>OBIEEBroker</servlet-name>
  <url-pattern>/obieebroker</url-pattern>
</servlet-mapping>

```

Following this step ensures that the OBIEEBroker servlet will be used to access your application from Oracle Business Intelligence.

10. Create an EAR deployment profile for the application, as follows:
 - a. Right-click the global application and select **Application Properties**.
 - b. From the left pane, select **Deployment**, then click **New** on the right pane to create a new deployment profile.
 - c. For **Archive Type**, select **EAR File**. Then, provide a name for the deployment profile (for example, MyApplication).

The deployment profile name is used as the name displayed in the list of deployments in Oracle WebLogic Server.

- d. From the left pane, select **Application Assembly**. Then, on the right pane, select the **webapp** deployment profile of your Web Project.

Following this step ensures that the WAR file from your Web Project is included in the EAR file.

11. Under **Application Resources**, select **Descriptors > META-INF > weblogic-application.xml**.

12. On the left, select the Libraries tab.

13. Create two new Shared Library References, as follows:

- Create the first Shared Library Reference with the following characteristics:
 - **Library Name:** oracle.bi.integration.adf
 - **Implementation Version:** 11.1.1.2.0
- Create the second Shared Library Reference with the following characteristics:
 - **Library Name:** oracle.applcore.model
 - **Implementation Version:** 11.1.1.0.0

These two Shared Library References create the following entries in the weblogic-application.xml file for the application:

```

<library-ref>
  <library-name>oracle.bi.integration.adf</library-name>
  <implementation-version>11.1.1.2.0</implementation-version>
</library-ref>
<library-ref>
  <library-name>oracle.applcore.model</library-name>
  <implementation-version>11.1.1.0.0</implementation-version>
</library-ref>

```

14. Deploy the EAR file to Oracle WebLogic Server by right-clicking the global application, then selecting **Deploy > EAR_deployment_profile_name**. From the dialog that appears, select **Deploy to Application Server** and then follow the instructions in the wizard.

15. To verify that the application has been deployed, log in to the WebLogic Server Administration Console and click **Deployments** under **Your Deployed Resources**. Verify that your application appears in the list (for example, *obieebroker_app_name*).

Setting Up a JDBC Data Source in the WebLogic Server

You must set up a JDBC data source in Oracle WebLogic Server for your application.

To set up a JDBC data source in Oracle WebLogic Server:

1. Ensure that Oracle WebLogic Server is running. If it is not running, start it. For example, on Windows, run `MW_HOME\user_projects\domains\your_domain\bin\startWebLogic.cmd`.
2. Open the WebLogic Server Administration Console. For example, if your Oracle WebLogic Server is running locally on port 7001, go to `http://localhost:7001/console`.
3. Log in to the WebLogic Server Administration Console with the credentials you created when you set up your WebLogic domain.
4. On the Home Page, select **JDBC**, then select **Data Sources**.
5. Click **New**.
6. Provide information for your data source. For **Name** and **JNDI Name**, provide the DataSource Name you specified in the Application Module configuration for the application, without the `java:comp/env` context prefix (for example, `jdbc/ApplicationDBDS`). In addition, make sure to select the target on which you want to deploy the data source before exiting the wizard.
7. Click **Finish** when you are done providing JDBC data source settings.

Setting the Logging Level for the Deployed Application in Oracle WebLogic Server

The log file for the server to which your application is deployed (*server_name*-diagnostic.log) records information about your deployed application. You can find this file in the server-specific directory within your domain. For example, on Windows, the log file for the AdminServer is located in:

```
MW_HOME\user_projects\domains\your_domain\servers\AdminServer\logs
```

To set the logging level for your deployed application:

1. Open the Oracle WebLogic Server file `logging.xml` for editing. You can find this file in:

```
MW_HOME\user_projects\domains\your_domain\config\fmwconfig\servers\server_name
```

2. Within the `<loggers>` element, add the following child elements:

```
<logger name="oracle.bi.integration.adf" level="LOG_LEVEL"/>
<logger name="oracle.bi.integration.adf.v11g.obieebroker" level="LOG_LEVEL"/>
```

Log levels include SEVERE, WARNING, INFO, CONFIG, FINE, FINER, and FINEST. Refer to the Oracle WebLogic Server documentation for information about logger levels.

3. Save and close the file.
4. Restart Oracle WebLogic Server.

Importing Metadata from ADF Business Component Data Sources

You must complete the steps in "Setting Up ADF Business Component Data Sources" before you can import metadata from ADF Business Component data sources.

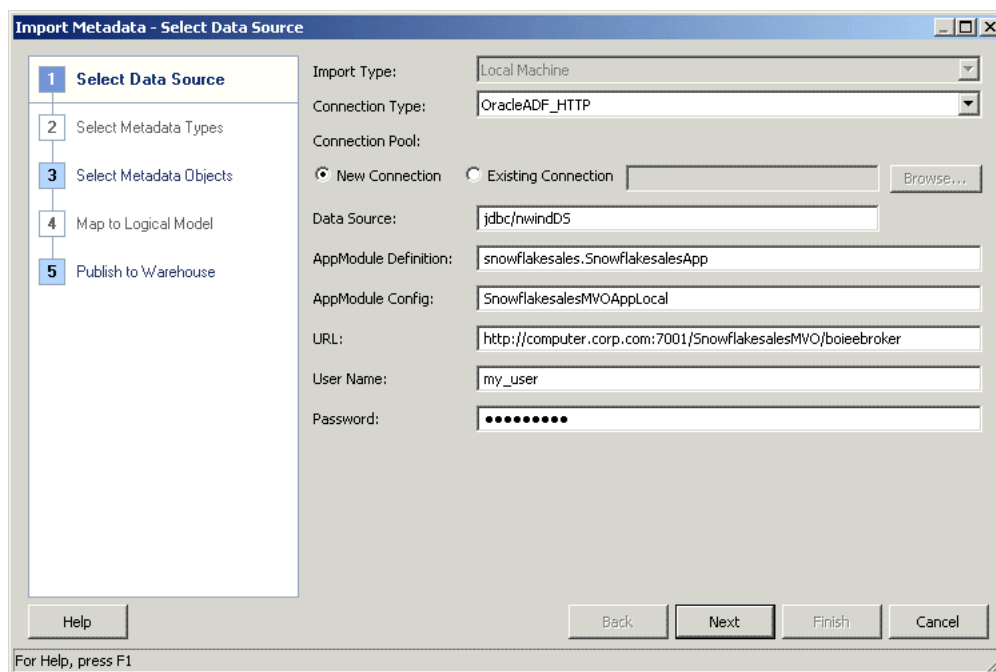
To import metadata from an ADF Business Component data source:

1. In the Administration Tool, select **File**, then select **Import Metadata**. The Import Metadata Wizard appears.

Note: If you have already defined an existing ADF Business Component data source and connection pool, you can right-click the connection pool in the Physical layer and select **Import Metadata**. The Import Metadata Wizard appears with the information on the Select Data Source screen pre-filled.

Figure 5–1 shows the Import Metadata Wizard.

Figure 5–1 Import Metadata Wizard: ADF Business Component Data Source



Note that the Map to Logical Model and Publish to Warehouse screens in the Import Metadata Wizard are reserved for a future release. Check the Oracle Business Intelligence chapter in *Oracle Fusion Middleware Release Notes* for updates about the functionality in these two screens.

2. In the Select Data Source screen, select **OracleADF_HTTP** for **Connection Type**. Then, provide the following values:
 - Select **New Connection**, or select **Existing Connection** if you already have a connection pool for this data source. Click **Browse** to locate and select an existing connection pool. If you select **Existing Connection**, you do not provide information for **Data Source**, **AppModule Definition**, **AppModule Config**, or **URL**, and the **User Name** and **Password** fields are pre-filled.

- Keep the **Data Source** field blank to use the default JDBC data source configured in the application module. You only need to provide data source information (a JDBC data source name, such as jdbc/nWindORA05) if you want to use a different data source than the one set up in the application module.
- For **AppModule Definition**, provide the fully qualified Java package name of the Root Application Module to which you want to connect, such as `oracle.apps.fii.receivables.model.RootAppModule`, or `snowflakesales.SnowflakeSalesApp`.
- For **AppModule Config**, provide the name of the configuration you want to use in your connection, such as `RootAppModuleShared` or `SnowflakeSalesAppLocal`. See step 3 of "[Deploying the Application EAR File to Oracle WebLogic Server from JDeveloper](#)" for more information.
- For **URL**, provide the URL to the Oracle Business Intelligence broker servlet, in the format:

```
http://host:port/APP_DEPLOYMENT_NAME/obieebroker
```

For example:

```
http://localhost:7001/MyApp/obieebroker
```

The URL is case-sensitive.

- For **User Name** and **Password**, provide a valid user name and password for the Oracle ADF application. The user name and password must be set up and authenticated in the Oracle WebLogic Server security realm.

When you have finished providing information in the Select Data Source screen, click **Next**. The Select Metadata Objects screen appears.

3. Select the objects you want to import in the **Available** list and move them to the **Selected** list, using the > (Import selected) or >> (Import all) buttons. You can also move objects from the **Selected** list back to the **Available** list, using the < (Remove selected) and << (Remove all) buttons.

To search for a particular item, enter a keyword in the **Find** box and then click **Find Down** or **Find Up**.

Select **Show complete structure** to view all objects, including those that have already been imported. Deselecting this option shows only the objects that are available for import. When this option is selected, objects that have already been imported appear grayed out.

When you move the items from the **Available** list to the **Selected** list, the Connection Pool dialog opens, showing the values that you provided in the Select Data Source screen of the Import Metadata Wizard. Optionally, click the Miscellaneous tab and provide the name of a SQL Bypass database in the **SQL Bypass Database** field. Then, click **OK**. If you do not want to specify a SQL Bypass database, click **Cancel**.

See "[About Specifying a SQL Bypass Database](#)" for more information.

4. Click **Finish**.
5. To validate that your import was successful, expand the database object for the ADF Business Component data source in the Physical layer. Then, right-click a physical table and click **View Data**. If the appropriate data is displayed, the import completed successfully.

Enabling the Ability to Pass Custom Parameters to the ADF Application

Some ADF applications have custom properties defined on the ApplicationModule, such as EFFECTIVE_DATE or TREE_VERSION. You can include these custom properties in your application queries, and the Oracle BI Server will pass them to the ADF application. To enable this feature, you must register the custom properties as a static repository variable using the Administration Tool.

You cannot use this feature to pass any custom property to your ADF application. Only certain custom properties, like EFFECTIVE_DATE and TREE_VERSION, are supported.

To register custom properties:

1. Open your repository in the Administration Tool.
2. Select **Manage**, then select **Variables**.
3. Select **Action > New > Repository > Variable**.
4. For **Name**, enter ADF_PARAM_LIST. Do not enter the name of the custom property as the name of the variable.
5. Ensure that the **Type** is **Static**.
6. For **Default Initializer**, enter the name or names of the custom properties as a character string. If you have multiple custom properties, include them as a comma-delimited list. For example:

```
'PARAM_EFFECTIVE_DATE'
```

```
'PARAM_EFFECTIVE_DATE, ApplicationIdBind, KeyFlexfieldCodeBind'
```

7. Click **OK**.
8. Save and close the repository.

After you register the custom properties as a repository variable, you can include these variables in queries. For example:

```
set variable PARAM_EFFECTIVE_DATE=2001-01-01 : SELECT c1 FROM t1;
```

or

```
set variable ApplicationIdBind = '0', KeyFlexfieldCodeBind = 'KFF1' :
select_physical ApplicationID, KeyFlexfieldCode, DataSecurityObjectName,
SegmentLabelCode from adfdb..."AppModule.KFFHierFilterV01";
```

Note that when you are including a custom property of type PARAM_EFFECTIVE_DATE, the date format for the property value must be in the format yyyy-mm-dd.

Propagating Labels and Tooltips from ADF Business Component Data Sources

You can propagate user interface hints, such as labels and tooltips, from ADF Business Component data sources to display when users work with analyses. When translated labels and tooltips (based on user locale) are maintained within an ADF Business Component data source, you can query the data source to access this translated data. You use the Administration Tool to configure presentation columns to use when creating analyses.

This section contains the following topics:

- [What are Labels and Tooltips?](#)

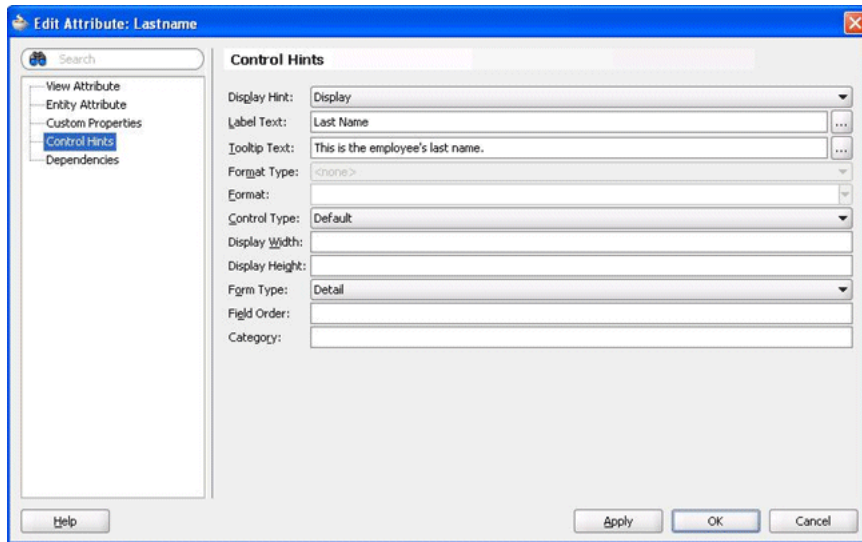
- [About the Session Variable Naming Scheme for UI Hints](#)
- [About Determining the Physical Column for a Presentation Column](#)
- [Initializing Session Variables Automatically for Propagating UI Hints](#)
- [Example of Using UI Hints From an Oracle ADF Data Source When Creating Analyses](#)
- [Using XML Code in Initialization Blocks to Query UI Hints](#)

What are Labels and Tooltips?

A label is the text that is used in prompts or table headers that precedes the value of a data item. A tooltip is the text that is displayed when a user hovers the mouse pointer over the item. Each attribute of a view object (VO) has an associated label and tooltip. A view object is the Oracle Application Development Framework component that enables a developer to work easily with SQL query results. The propagation of UI hints enables a presentation column in the Administration Tool to use a label and tooltip as its Custom display name and Description respectively.

Figure 5–2 shows the Label Text and Tooltip Text options in the Edit Attribute dialog in Oracle JDeveloper.

Figure 5–2 Edit Attribute Dialog in JDeveloper for Label and Tooltip Options



About the Session Variable Naming Scheme for UI Hints

Session variable names are generated by the Oracle BI Enterprise Edition broker servlet in Oracle WebLogic Server in the following format:

ADF_UI Hint Type_Database Name_VO's Name_Attribute's Name

Where:

UI Hint Type is either LABEL or TOOLTIP, depending on the UI hint type that the session variable represents.

Database Name is the value of the "database" attribute of the ADFQuery element in the XML query. Special characters such as single quotes ('), double quotes ("), and spaces are replaced by the underscore character.

VO's Name is the name of the View Object to which the attribute belongs. Oracle ADF prohibits special characters and spaces in the name.

Attribute's Name is the name of the attribute that the session variable represents. Oracle ADF prohibits special characters and spaces in the name.

Every character in the session variable name is uppercase. For example, the XML query in [Example 5-3](#) generates four session variables with the following names:

```
ADF_LABEL_MY_ORCLADF_EMPLOYEESVIEW_FIRSTNAME
ADF_TOOLTIP_MY_ORCLADF_EMPLOYEESVIEW_FIRSTNAME
ADF_LABEL_MY_ORCLADF_EMPLOYEESVIEW_LASTNAME
ADF_TOOLTIP_MY_ORCLADF_EMPLOYEESVIEW_LASTNAME
```

About Determining the Physical Column for a Presentation Column

As required by the naming scheme for session variables, each presentation column must map to a physical column. When you select **Externalize Display Names > Generate ADF Label** or **Externalize Descriptions > Generate ADF Tooltip** for a presentation layer object, then the physical column is located using the following rules:

1. Examine the presentation column and determine its logical column. If the logical column is derived from an existing logical column, then the physical column cannot be found.
2. If the default aggregation rule for the logical column is not None or Sum, then the physical column cannot be found. It does not make sense semantically to use the ADF UI hints for aggregation rules other than Sum.
3. A logical column can be mapped to physical columns by multiple logical table sources. Only logical table sources that are not disabled are searched.
4. Do not search logical table sources that map the logical column using non-trivial expressions (that is, anything more than a physical column name). If no logical table sources are searched, then the physical column cannot be found.
5. From the remaining ordered list of logical table sources, examine the physical column that is mapped by the first logical table source. The physical column must be mapped to a VO attribute. In other words the physical column must be part of a physical database of type OracleADF11g.
 - If this condition is satisfied, then the physical column for obtaining UI hints is found.
 - If this condition is not satisfied, then continue to examine the physical column that is mapped by the next logical table source until the physical column that is mapped to a VO attribute is found.

If all logical table source are searched without satisfying the condition, then the physical column cannot be found.

If the physical column for obtaining UI hints is found using these rules, then the custom display name or description is populated with a session variable that has a name based on a predetermined naming scheme. See ["About the Session Variable Naming Scheme for UI Hints"](#) for more information.

If the physical column for obtaining UI hints is not found using these rules, then the **Generate ADF Label** and **Generate ADF Tooltip** options are shown as disabled in the right-click menu.

As an alternative to using the physical column found using these rules, you can use XML code in an initialization block to initialize your own session variables with ADF UI hints. You must then enter these session variable names in the **Custom display name** and **Custom description** fields manually. See ["Using XML Code in Initialization Blocks to Query UI Hints"](#) for more information.

Initializing Session Variables Automatically for Propagating UI Hints

If the **Externalize Display Names > Generate ADF Label** and **Externalize Descriptions > Generate ADF Tooltip** options were used to successfully generate the session variable names for UI hints from Oracle ADF, then the session variables are created and initialized when Oracle BI Presentation Services queries them during the session. The variables are not created and initialized during the session logon stage for performance reasons. Instead, the variables are created and initialized when they are needed by a specific query within a session, using the **Allow deferred execution** feature.

When Presentation Services queries the custom display names and custom descriptions through ODBC, the Oracle BI Server checks if the associated session variables have been created. If they have not been created, then the Oracle BI Server dynamically generates the appropriate XML query (as described in ["Using XML Code in Initialization Blocks to Query UI Hints"](#)) to query the UI hints from the Oracle ADF data source. The Oracle BI Server uses the UI hints to create and initialize the session variables. As an optimization, the Oracle BI Server queries UI hints per VO; that is, if the Oracle BI Server needs the UI hints of a VO's attributes, then the UI hints for all the attributes under the VO are queried and propagated through session variables.

Example of Using UI Hints From an Oracle ADF Data Source When Creating Analyses

The following example shows how you can use UI hints from an Oracle ADF data source when creating analyses.

The following prerequisites must be met:

- UI hints must have been configured in the Oracle ADF data source.
- A working repository must have been configured for the Oracle ADF data source in the Administration Tool.

To use UI hints from an Oracle ADF data source when creating analyses:

1. Suppose that the repository contains a presentation column named "LastName." On the General tab of the Presentation Column dialog, the **Custom display name** and **Custom description** fields are not selected.

Right-click the column in the Presentation layer and select first **Externalize Display Names > Generate ADF Label**, then **Externalize Descriptions > Generate ADF Tooltip** to generate the strings that populate the **Custom display name** and **Custom description** fields.

You can also use these options from the right-click menu of a presentation table to generate the strings for all the columns in that table.

2. View the UI hints:
 - a. Sign in to Oracle Business Intelligence.
 - b. Create a new analysis using the subject area for which you obtained UI hints.
 - c. In the Subject Areas pane, expand the Employee folder to see the UI hints that have been propagated from the Oracle ADF data source.

The LastName column displays as "Last Name" (the label value from the Oracle ADF data source). When you hover the mouse pointer over the column, the tip displays as "This is the employee's last name" (the tooltip value from the Oracle ADF data source).

For information about creating analyses, see *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition*.

Using XML Code in Initialization Blocks to Query UI Hints

As an alternative to using the automated system described in the previous section, you can use specialized XML code in place of SQL statements in initialization blocks to query the data source for UI hints, within a single repository and subject area. You use the ADFQuery element, which has three attributes that are named mode, database, and locale. The element requires zero or more child elements. The syntax of the element is as follows:

```
<?xml version="1.0" encoding="iso-8859-1" standalone="yes"?>
<ADFQuery mode="{Mode}" database="{Database Name}"
locale="VALUEOF(NQ_SESSION.WEBLANGUAGE)">
  <ViewObject><![CDATA[{VO Name}]]></ViewObject>
  <Attribute>
  <ViewObject><![CDATA[{Attribute VO Name}]]></ViewObject>
  <Name><![CDATA[{Attribute Name}]]></Name>
  </Attribute>
</ADFQuery>
```

where

{Mode} specifies what you want to query:

- label for querying attributes' label
- tooltip for querying attributes' tooltip
- ui_hints for querying attributes' label and tooltip

{Database Name}

Use the name of the physical database object in the Administration Tool, which contains the physical columns that correspond to the attributes in the Oracle ADF data source.

{VO Name}

Use the name of the View Object to obtain the UI hints of all attributes in it.

{Attribute VO Name}

Use the name of the View Object that contains the attribute.

{Attribute Name}

Use the name of the attribute that belongs to the associated View Object to obtain the UI hints of this attribute.

Example 5-1 Querying Labels for All View Objects

No child elements must be included in the ADFQuery element, if the UI hints of all attributes in all View Objects are queried. For example, to query the labels of all attributes in all View Objects under the My_orclADF physical database object, use the following XML code:

```
<?xml version="1.0" encoding="iso-8859-1" standalone="yes"?>
```



```
<ADFQuery mode="label" database="My_orclADF"
locale="VALUEOF (NQ_SESSION.WEBLANGUAGE)" ">
</ADFQuery>
```

Example 5–2 Querying Tooltips for Specific View Objects

The ADFQuery element can contain zero or more child elements named ViewObject if UI hints of all attributes in specific View Objects are queried. Each ViewObject element has a text content that contains the View Object's name. The ViewObject element is used to specify the View Objects from which the UI hints of all attributes are queried. For example, to query the tooltips of all attributes in the View Object that is named EmployeesView and CustomersView under the My_orclADF physical database object, use the following XML code:

```
<?xml version="1.0" encoding="iso-8859-1" standalone="yes"?>
<ADFQuery mode="tooltip" database="My_orclADF"
locale="VALUEOF (NQ_SESSION.WEBLANGUAGE)" ">
  <ViewObject><![CDATA[EmployeesView]]></ViewObject>
  <ViewObject><![CDATA[CustomersView]]></ViewObject>
</ADFQuery>
```

Example 5–3 Querying UI Hints for Specific Attributes

The ADFQuery element can contain zero or more child elements named Attribute. Each Attribute element has two required child elements named ViewObject and Name. The Attribute element is used to specify the attributes from which the UI hints are queried. The ViewObject child element has a text content that contains the View Object's name. This element specifies the View Object that the attribute belongs to. The Name child element has a text content which contains the attribute's name. For example, to query the labels and tooltips of the attributes named Firstname and Lastname in the EmployeesView View Object under the My_orclADF physical database object, use the following XML code:

```
<?xml version="1.0" encoding="iso-8859-1" standalone="yes"?>
<ADFQuery mode="ui_hints" database="My_orclADF"
locale="VALUEOF (NQ_SESSION.WEBLANGUAGE)" ">
  <Attribute>
    <ViewObject><![CDATA[EmployeesView]]></ViewObject>
    <Name><![CDATA[Firstname]]></Name>
  </Attribute>
  <Attribute>
    <ViewObject><![CDATA[EmployeesView]]></ViewObject>
    <Name><![CDATA[Lastname]]></Name>
  </Attribute>
</ADFQuery>
```

After configuring the initialization blocks, you must manually enter the session variable names in the **Custom display name** and **Custom description** text fields for the appropriate presentation column.

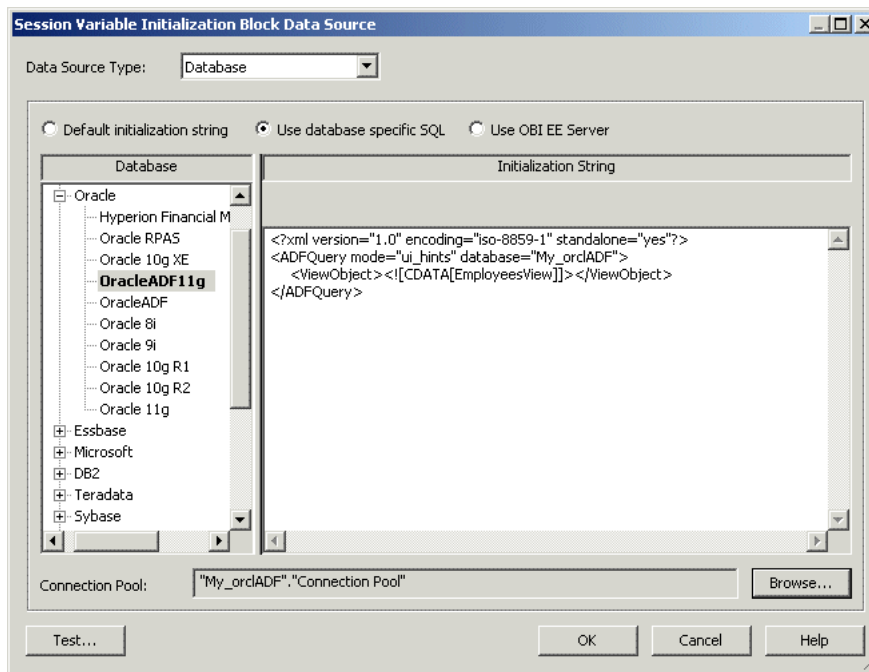
Follow the procedure in the example in ["Example of Using UI Hints From an Oracle ADF Data Source When Creating Analyses"](#), but replace the first step with the following ones:

1. Create session initialization blocks in the Administration Tool.
 - a. In the Session Variable Initialization Block Data Source dialog, enter the Initialization string.

In this example, the initialization block queries both the label and tooltip of all attributes in a View Object named EmployeesView. [Figure 5–3](#) shows the setup

of a session variable initialization block with an appropriate Oracle ADF UI hint query. "My_orclADF"."Connection Pool" is a connection pool for an Oracle ADF data source.

Figure 5–3 Setting Up a Session Variable Initialization Block Data Source with an Oracle ADF UI Hints Query



- b. In the Session Variable Initialization Block dialog, select **Row-wise initialization** as the Variable Target.
- c. Click **Test** to test the query against the Oracle ADF data source.

In the results window, the first column contains the session variable names that are generated using the naming scheme. The second column contains the label and tooltip values from the Oracle ADF data source.

See ["About the Session Variable Naming Scheme for UI Hints"](#) for a description of the naming scheme.

2. Configure a custom display name and a description in presentation columns.


To find the presentation tables that can use the UI hints from the EmployeesView View Object, this example uses the Query Repository feature in the Administration Tool.

- a. Right-click a physical table (for example, EmployeesView), then select **Query Related Objects > Presentation > Presentation Table** from the menu.

The Query Related Objects dialog displays all the related presentation tables.

This example sets up a custom display name and custom description for columns in the Employee presentation table.

Figure 5–4 Using the Query Related Objects Feature to Find the Related Presentation Tables



- b. Select the required presentation table and click **Go To**.
This displays the selected presentation table.
- c. Expand the presentation table to view the presentation columns.
- d. Double-click the LastName presentation column to display the Presentation Column dialog.
- e. Select **Custom display name** and enter a value such as the following one:
`VALUEOF (NQ_SESSION.ADF_LABEL_MY_ORCLADF_EMPLOYEESVIEW_LASTNAME)`
- f. Select **Custom description** and enter a value such as the following one:
`VALUEOF (NQ_SESSION.ADF_TOOLTIP_MY_ORCLADF_EMPLOYEESVIEW_LASTNAME)`
- g. Click **OK**.
- h. Save the changes in the repository and restart the Oracle BI Server.

Setting Up Database Objects and Connection Pools

Properties for database objects and connection pools are typically set automatically when you import metadata from your data sources. However, in some cases you may want to adjust database or connection pool settings, or create a database object or connection pool manually. This chapter describes the properties of the database and connection pool objects in the Physical layer.

This chapter contains the following sections:

- [Setting Up Database Objects](#)
- [About Connection Pools](#)
- [Creating or Changing Connection Pools](#)
- [Setting Up Persist Connection Pools](#)

Setting Up Database Objects

Importing metadata from a data source automatically creates a database object for the schema, but you may need to adjust the database properties.

See "[System Requirements and Certification](#)" for information about supported data sources.

The following sections provide information about how to create or edit database objects in the Physical layer:

- [About Database Types in the Physical Layer](#)
- [Creating a Database Object Manually in the Physical Layer](#)
- [Specifying SQL Features Supported by a Data Source](#)

About Database Types in the Physical Layer

If you import the physical schema into the Physical layer, the database type is usually assigned automatically. The following list contains additional information about automatic assignment of database types:

- **Relational data sources.** During the import process, some ODBC drivers provide the Oracle BI Server with the database type. However, if the server cannot determine the database type, an approximate ODBC type is assigned to the database object. Replace the ODBC type with the closest matching entry from the **Database** list.

- **Multidimensional data sources.** Microsoft Analysis Services and SAP/BW are the only supported XMLA-compliant data sources currently available. After you import metadata from a multidimensional data source, check the database object and update the appropriate database type and version if necessary.

Creating a Database Object Manually in the Physical Layer

If you create a database object manually, you need to manually set up all database elements such as the connection pool, tables, and columns.

For multidimensional data sources, if you create the physical schema in the Physical layer of the repository, you need to create one database in the physical layer for each cube, or set of cubes, that belong to the same catalog (database) in the data source. A physical database can have more than one cube. However, all of these cubes must be in the same catalog in the data source.

Caution: It is strongly recommended that you import your physical schema.

To create a database object:

1. In the Administration Tool, in the Physical layer, right-click and select **New Database**.
Make sure that no object is selected when you right-click.
2. In the Database dialog, in the General tab, complete the fields using [Table 6–1](#) as a guide.

Table 6–1 Options in the General Tab of the Database Dialog

Option	Description
Data source definition: Database	The database type for your database. See " Specifying SQL Features Supported by a Data Source " for more information about using the Features tab to examine the SQL features supported by the specified database type.
Data source definition: CRM metadata tables	This property is only available for relational data sources and is for legacy Siebel Systems sources only. When selected, indicates that the definition of physical tables and columns for Siebel CRM tables was derived from the Siebel metadata dictionary.
Data source definition: Virtual Private Database	Identifies the physical database source as a virtual private database (VPD). When a VPD is used, returned data results are contingent on the user's authorization credentials. Therefore, it is important to identify these sources. These data results affect the validity of the query result set that is used with caching. See " Managing Performance Tuning and Query Caching " in <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i> . Note: If you select this option, you also should select the Security Sensitive option in the Session Variable dialog. See " Creating Session Variables " for more information.
Persist connection pool	To use a persistent connection pool, you must set up a temporary table first. See " Setting Up Persist Connection Pools " for more information.

Table 6–1 (Cont.) Options in the General Tab of the Database Dialog

Option	Description
Allow populate queries by default	When selected, allows everyone to execute <code>POPULATE SQL</code> . If you want most, but not all, users to be able to execute <code>POPULATE SQL</code> , select this option and then limit queries for specific users or groups. See "Setting Query Limits" for more information.
Allow direct database requests by default	<p>When selected, allows all users to execute physical queries. The Oracle BI Server sends unprocessed, user-entered, physical SQL directly to an underlying database. The returned results set can be rendered in Oracle BI Presentation Services, and then charted, rendered in a dashboard, and treated as an Oracle BI request.</p> <p>If you want most, but not all, users to be able to execute physical queries, select this option and then limit queries for specific users or groups. See "Setting Query Limits" for more information.</p> <p>Caution: If configured incorrectly, this option can expose sensitive data to an unintended audience. See "When to Allow Direct Database Requests by Default" for more information.</p> <p>For more information about executing physical SQL, see <i>"Working with Direct Database Requests"</i> in <i>Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition</i>.</p>

When to Allow Direct Database Requests by Default

The property **Allow direct database requests by default** lets all users execute physical queries. If configured incorrectly, it can expose sensitive data to an unintended audience. Use the following recommended guidelines when setting this database property:

- The Oracle BI Server should be configured to accept connection requests only from a computer on which the Oracle BI Server, Oracle BI Presentation Services, or Oracle BI Scheduler are running. This restriction should be established at the TCP/IP level using the Oracle BI Presentation Services IP address. This allows only a TCP/IP connection from the IP address of Oracle BI Presentation Services.
- To prevent users from running `nqcmd` (a utility that executes SQL scripts) by logging in remotely to this computer, you should disallow access by the following to the computer on which you installed Oracle BI Presentation Services:
 - TELNET
 - Remote shells
 - Remote desktops
 - Teleconferencing software (such as Windows NetMeeting)

If necessary, you might want to make an exception for users with administrator permissions.

- Only users with administrator permissions should be allowed to perform the following tasks:
 - TELNET into the Oracle BI Server and Oracle BI Presentation Services computers to perform tasks such as running `nqcmd` for cache seeding.

- Access the advanced SQL page of Answers to create requests. For more information, see *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition*.
- Set up group/user-based permissions on Oracle BI Presentation Services to control access to editing (preconfigured to allow access by Oracle BI Presentation Services administrators) and executing (preconfigured to not allow access by anyone) direct database requests. For more information, see *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

Specifying SQL Features Supported by a Data Source

When you import metadata or specify a database type in the General tab of the Database dialog, the set of features for that database object is automatically populated with default values appropriate for the database type. These are the SQL features that the Oracle BI Server uses with this data source.

When a feature is marked as supported (checked) in the Features tab of the Database dialog, the Oracle BI Server typically pushes the function or calculation down to the data source for improved performance. When a function or feature is not supported in the data source, the calculation or processing is performed in the Oracle BI Server.

The supported features list in the Features tab uses the feature defaults defined in the DBFeatures.INI file, located in `ORACLE_INSTANCE\config\OracleBIServerComponent\coreapplication_obisn`. Although you should not modify this file directly, it can be useful to look at this file to compare the features supported by different data source types.

You can tailor the query features for a data source. For example, a new version of a data source may be released with updated feature support that is not reflected in the Oracle BI Server defaults. In this case, you can update the settings in the Features tab to reflect the actual features supported by the new version of the data source. Or, if a data source supports a particular feature (such as left outer join queries) but you want to prohibit the Oracle BI Server from sending such queries to a particular data source, you can change this default setting in the Features tab. A third situation is when you have federated data sources that execute functions differently. To ensure query results are consistent, you can disable the appropriate functions so that the calculations are performed in a consistent manner in the Oracle BI Server.

Caution: Be very careful when modifying the set of supported features in the Features tab. If you enable SQL features that the data source does not support, your query may return errors and unexpected results. If you disable supported SQL features, the server could issue less efficient SQL to the data source.

In most cases, you should keep the default values. If you do change the defaults to mark a feature as supported in the Features tab, make sure that the feature is actually supported by the data source.

To specify SQL features supported by a data source:

1. In the Administration Tool, in the Physical layer, double-click the database for which you want to specify SQL features.
2. In the Database dialog, click the Features tab.
3. In the Features tab, use the information in [Table 6–2](#) to help you specify properties for each SQL feature.

Table 6–2 Options in the Features Tab of the Database Dialog

Option	Description
Feature	The name of the database feature, such as COUNT_DISTINCT_SUPPORTED.
Value	Shows the current value for the given feature. Selected indicates that the feature is supported in the data source, and that the function or feature should be performed in the data source rather than in the Oracle BI Server. Some features show a default value in the Value column rather than selected/not selected, such as 10 for MAX_ENTRIES_PER_IN_LIST. It is strongly recommended that you keep the default selections and default values.
Default	Shows the default value for the given feature. The defaults listed in this column are specified in the file DBFeatures.INI.
Find	Lets you type a string to help you locate a feature in the list.
Find Again	This option becomes available after you click Find . It lets you perform multiple searches for the same string.
Query DBMS	This button is only used if you are installing and querying a data source that has no set of feature defaults in the Oracle BI Server. It lets you query this type of data source for Feature table entries so that you can find out which SQL features it supports. You can then change the entries that appear in the Features tab based on your query results. This button is not available if you are using an XML or a multidimensional data source. Caution: Be very careful when using the Query DBMS feature. The results of the features query are not always an accurate reflection of the SQL features actually supported by the data source. When using this feature, you should verify that the list of supported features in the Features tab matches the actual features supported by your data source. Refer to the documentation for your data source for details.
Reset to defaults	This button restores the default values for this data source type from the file DBFeatures.INI.

Note: Do not change the OPTIMIZE_MDX_FILTER_QUALIFICATION feature. This parameter is reserved for a future release.

About Connection Pools

The connection pool is an object in the Physical layer that describes access to the data source. It contains information about the connection between the Oracle BI Server and that data source.

The Physical layer in the Administration Tool contains at least one connection pool for each database. When you create the Physical layer by importing a schema for a data source, the connection pool is created automatically. You can configure multiple connection pools for a database. Connection pools allow multiple concurrent data source requests (queries) to share a single database connection, reducing the overhead of connecting to a database.

Note: It is recommended that you create a dedicated connection pool for initialization blocks. See "[About Connection Pools for Initialization Blocks](#)" for more information.

For each connection pool, you must specify the maximum number of concurrent connections allowed. After this limit is reached, the connection request waits until a connection becomes available.

Increasing the allowed number of concurrent connections can potentially increase the load on the underlying database accessed by the connection pool. Test and consult with your DBA to make sure the data source can handle the number of connections specified in the connection pool. Also, if the data sources have a charge back system based on the number of connections, you might want to limit the number of concurrent connections to keep the charge-back costs down.

In addition to the potential load and costs associated with the database resources, the Oracle BI Server allocates shared memory for each connection upon server startup. This raises the number of connections and increases Oracle BI Server memory usage.

About Connection Pools for Initialization Blocks

It is recommended that you create a dedicated connection pool for initialization blocks. This connection pool should not be used for queries.

Additionally, it is recommended that you isolate the connections pools for different types of initialization blocks. This also makes sure that authentication and login-specific initialization blocks do not slow down the login process. The following types should have separate connection pools:

- All authentication and login-specific initialization blocks such as language, externalized strings, and group assignments.
- All initialization blocks that set session variables.
- All initialization blocks that set repository variables. These initialization blocks should always be run using credentials with administrator privileges.

Be aware of the number of these initialization blocks, their scheduled refresh rate, and when they are scheduled to run. Typically, it would take an extreme case for this scenario to affect resources. For example, refresh rates set in minutes, greater than 15 initialization blocks that refresh concurrently, and a situation in which either of these scenarios could occur during prime user access time frames.

Initialization blocks should be designed so that the maximum number of Oracle BI Server variables may be assigned by each block. For example, if you have five variables, it is more efficient and less resource intensive to construct a single initialization block containing all five variables. When using one initialization block, the values are resolved with one call to the back end tables using the initialization string. Constructing five initialization blocks, one for each variable, would result in five calls to the back end tables for assignment.

See "[Working with Initialization Blocks](#)" for more information about these objects.

Creating or Changing Connection Pools

Typically, database objects and connection pools are created automatically when you import physical schemas, for both relational and multidimensional data sources. If you did not import physical schemas, you must create a database object before you create a

connection pool. You create or change a connection pool in the Physical layer of the Administration Tool.

If you have already defined an existing database and connection pool, you can right-click the connection pool in the Physical layer and select **Import Metadata** to import metadata for this data source. The Import Metadata Wizard appears with the information on the Select Data Source screen pre-filled. See [Chapter 4, "Importing Metadata and Working with Data Sources"](#) for information about the Import Wizard.

To automate connection pool changes for use in a process such as production migration, consider using the Oracle BI Server XML API. See "About the Oracle BI Server XML API" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

To create or change a connection pool:

1. In the Physical layer of the Administration Tool, right-click a database and select **New Object**, then select **Connection Pool**. Or, double-click an existing connection pool.
2. Specify or adjust the properties as needed, then click **OK**.

The following sections describe how to set properties in the various tabs of the Connection Pool dialog:

- [Setting Connection Pool Properties in the General Tab](#)
- [Setting Connection Pool Properties in the Connection Scripts Tab](#)
- [Setting Connection Pool Properties in the XML Tab](#)
- [Setting Connection Pool Properties in the Write Back Tab](#)
- [Setting Connection Pool Properties in the Miscellaneous Tab](#)

Setting Connection Pool Properties in the General Tab

This section describes the properties in the General tab of the Connection Pool dialog. The General tab is available for all data sources.

To set general properties for connection pools:

- In the Connection Pool dialog, click the General tab, and then complete the fields using the information in [Table 6-3](#) and [Table 6-4](#).

The properties listed in the General tab vary according to the data source type. For example, XMLA data sources have a connection pool property for **URL**, while relational and XML data sources have the option **Require fully qualified table names**.

This section contains the following topics:

- [Common Connection Pool Properties in the General Tab](#)
- [Multidimensional Connection Pool Properties in the General Tab](#)

Common Connection Pool Properties in the General Tab

This section describes connection pool properties in the General tab that are common among most data source types.

[Figure 6-1](#) shows the General tab of the Connection Pool dialog, for an OCI data source.

Figure 6–1 General Tab of the Connection Pool Dialog: OCI Data Source




Table 6–3 describes the properties in the General tab of the Connection Pool dialog that are common for different data source types.

Table 6–3 Connection Pool Properties in the General Tab: Common Properties

Property	Description
Name	The name for the connection pool. A name is assigned automatically for connection pools created upon import.
Permissions	Use this option to assign permissions for individual users or application roles to access the connection pool. For example, you can set up a privileged group of users to have its own connection pool. This feature is not intended to be used for data access security. For example, connection pool permissions do not protect cache entries. Refer to Chapter 13 for complete information on data access security in Oracle Business Intelligence.
Call interface	Identifies the application programming interface (API) with which to access the data source. Some databases can be accessed using native APIs, some use ODBC, and some work both ways. If the call interface is XML, the XML tab is available but options that do not apply to XML data sources are not available.

Table 6–3 (Cont.) Connection Pool Properties in the General Tab: Common Properties

Property	Description
Maximum connections	<p>The maximum number of connections allowed for this connection pool. The default is 10. This value should be determined by the database make and model and the configuration of the hardware for the computer on which the database runs, as well as the number of concurrent users who require access.</p> <p>For Microsoft Analysis Services data sources, you might encounter 503 Service Not Available errors if the Max Connections setting in the connection pool (default 10) is greater than the XMLA MaxThreadsPerClient setting configured in Analysis Services (default 4). To avoid these errors, increase the MaxThreadsPerClient setting in the msmdpump.ini file, or reduce the Max Connections setting in the repository connection pool.</p> <p>See also "Avoiding Query Failures Due to Limited Connections with TimesTen Data Sources" for related information.</p> <p>Note: For deployments with Oracle BI Interactive Dashboards pages, consider estimating this value at 10% to 20% of the number of simultaneous users multiplied by the number of requests on a dashboard. This number can be adjusted based on usage. The total number of all connections in the repository should be less than 800. To estimate the maximum connections needed for a connection pool dedicated to an initialization block, you might use the number of users concurrently logged on during initialization block execution.</p>

Table 6–3 (Cont.) Connection Pool Properties in the General Tab: Common Properties

Property	Description
Require fully qualified table names	<p>Select this option if the database or database configuration requires fully qualified table names. This option is not available for some data source types.</p> <p>When this option is selected, all requests sent from the connection pool use fully qualified names to query the underlying database. The fully qualified names are based on the physical object names in the repository. If you are querying the same tables from which the Physical layer metadata was imported, you can safely select this option. If you have migrated your repository from one physical database to another physical database that has different database and schema names, the fully qualified names would be invalid in the newly migrated database. In this case, if you do not select this option, the queries will succeed against the new database objects.</p> <p>For some data sources, fully qualified names might be safer because they guarantee that the queries are directed to the desired tables in the desired database. For example, if the RDBMS supports a master database concept, a query against a table named Customer first looks for that table in the master database, and then looks for it in the specified database. If the table named Customer exists in the master database, that table is queried, not the table named Customer in the specified database.</p> <p>It is sometimes necessary to select this option when you are using an Oracle Database, and you are accessing the database with a user that is not the owner of the schema containing the tables. When the Oracle Database interprets table names in SQL, it assumes that the user that made the query is the owner if the table name is not fully qualified in the query. This can result in an incorrect qualified name.</p> <p>For example, if the user SAMPLE creates a table called CUSTOMER, the fully qualified table name is SAMPLE.CUSTOMER. When the SAMPLE user references the CUSTOMER table in a query, the Oracle Database assumes the fully qualified table name is SAMPLE.CUSTOMER, and the access is successful. However, if the JANEDOE user references the CUSTOMER table in a query, the Oracle Database assumes the fully qualified table name is JANEDOE.CUSTOMER, and a "Table or view not found" error can result. To enable access for JANEDOE, you must select Require fully qualified table names in the connection pool so that the Oracle BI Server specifies SAMPLE.CUSTOMER in all queries.</p>
Data source name	<p>The name of the data source to which you want this connection pool to connect and send physical queries. The value you enter in this field depends on the selected call interface:</p> <ul style="list-style-type: none"> ■ If the call interface is OCI, enter a full connect string or a net service name from the tnsnames.ora file you set up within the Oracle Business Intelligence environment, in <i>ORACLE_HOME/network/admin</i>. ■ If you are using a native interface for a different database, enter the name of the database for that system. ■ If the call interface is ODBC, the data source name field displays a list containing all the User and System DSNs defined for ODBC on the local computer. Select the correct one for the data source to which you want connect.

Table 6–3 (Cont.) Connection Pool Properties in the General Tab: Common Properties

Property	Description
Shared logon	<p>Select this option if you want all users whose queries use the connection pool to access the underlying database using the same user name and password.</p> <p>If this option is selected, then all connections to the database that use the connection pool use the user name and password specified in the connection pool, even if the user has specified a database user name and password in the DSN (or in user configuration).</p> <p>If this option is not selected, connections through the connection pool use the database user ID and password specified in the DSN or in the user profile.</p>
Enable connection pooling	<p>When selected, allows a single database connection to remain open for the specified time for use by future query requests. Connection pooling saves the overhead of opening and closing a new connection for every query. If you do not select this option, each query sent to the database opens a new connection.</p>
Timeout	<p>Specify the amount of time and in what increment (such as minutes) that a connection to the data source remains open after a request completes. During this time, new requests use this connection rather than open a new one (up to the number specified for the maximum connections). The time is reset after each completed connection request.</p>
Use multithreaded connections	<p>When this option is selected, the Oracle BI Server terminates idle physical queries (threads). When not selected, one thread is tied to one database connection (number of threads = maximum connections). Even if threads are idle, they consume memory.</p> <p>The parameter <code>DB_GATEWAY_THREAD_RANGE</code> in the Server section of <code>NQSConfig.ini</code> establishes when the Oracle BI Server terminates idle threads. The lower number in the range is the number of threads that are kept open before the Oracle BI Server takes action. If the number of open threads exceeds the low point in the range, the Oracle BI Server terminates idle threads. For example, if <code>DB_GATEWAY_THREAD_RANGE</code> is set to 40-200 and 75 threads are open, the Oracle BI Server terminates any idle threads.</p>
Parameters supported	<p>If this option is not selected, and the database features table supports parameters, special code executes that allows the Oracle BI Server to push filters (or calculations) with parameters to the database. The Oracle BI Server does this by simulating parameter support within the gateway/adaptor layer by sending extra <code>SQLPrepare</code> calls to the database.</p>

Table 6–3 (Cont.) Connection Pool Properties in the General Tab: Common Properties

Property	Description
Isolation level	<p>For ODBC and DB2 gateways only. The value sets the transaction isolation level on each connection to the back-end database. The isolation level setting controls the default transaction locking behavior for all statements issued by a connection. Only one option can be set at a time. It remains set for that connection until it is explicitly changed.</p> <p>The following options are available:</p> <p>Dirty read. Implements dirty read (isolation level 0 locking). This is the least restrictive isolation level. When this option is set, it is possible to read uncommitted or dirty data, change values in the data, and have rows appear or disappear in the data set before the end of the transaction.</p> <p><i>Dirty data</i> is data that needs to be cleaned before being queried to obtain correct results (for example, duplicate records, records with inconsistent naming conventions, or records with incompatible data types).</p> <p>Committed read. Specifies that shared locks are held while the data is read to avoid dirty reads. However, the data can be changed before the end of the transaction, resulting in non-repeatable reads or phantom data.</p> <p>Repeatable read. Places locks on all data that is used in a query, preventing other users from updating the data. However, new phantom rows can be inserted into the data set by another user and are included in later reads in the current transaction.</p> <p>Serializable. Places a range lock on the data set, preventing other users from updating or inserting rows into the data set until the transaction is complete. This is the most restrictive of the four isolation levels. Because concurrency is lower, use this option only if necessary.</p>

Multidimensional Connection Pool Properties in the General Tab

This section describes connection pool properties in the General tab that are specific to multidimensional data sources.

Figure 6–2 shows the General tab of the Connection Pool dialog, for an Essbase data source.

Figure 6–2 General Tab of the Connection Pool Dialog: Essbase Data Source

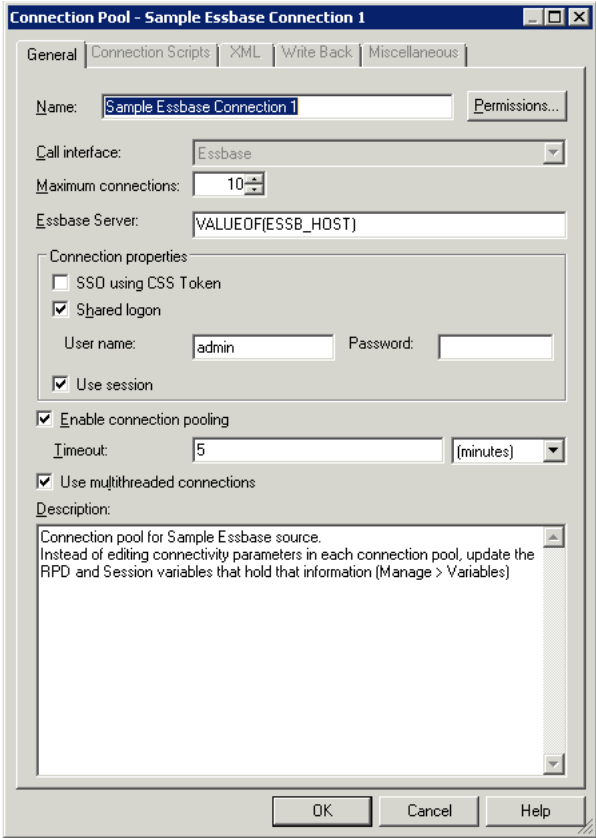


Table 6–4 describes the properties in the General tab of the Connection Pool dialog that are specific to multidimensional data sources. Note that some properties only appear for certain types of multidimensional data sources.

Table 6–4 Connection Pool Properties in the General Tab: Multidimensional Data Source Properties

Property	Description
URL	This property is only displayed for XMLA data sources. Specify the URL to connect to the XMLA provider. This URL points to the XMLA virtual directory of the computer hosting the cube. This virtual directory must be associated with msxisapi.dll (part of the Microsoft XML for Analysis SDK installation). For example, the URL might look like the following: <code>http://SDCDL360i101/xmla/msxisap.dll</code>
Essbase Server	This property is only displayed for Essbase data sources. Specify the name of the Essbase server.
SSO using CSS Token	This property is only displayed for Essbase and Hyperion Financial Management data sources. Select this option if you want to authenticate against Essbase or Hyperion Financial Management using a shared token, rather than using a set of shared credentials in the connection pool. Make sure to select either the SSO using CSS Token or Shared logon options when creating a connection pool for Essbase or Hyperion Financial Management. Do not leave both options blank.

Table 6–4 (Cont.) Connection Pool Properties in the General Tab: Multidimensional Data Source Properties

Property	Description
Data Source Information: Data Source	Specify the vendor-specific information used to connect to the multidimensional data source. Consult your multidimensional data source administrator for setup instructions because specifications can change. For example, if you use v1.0 of the XML for Analysis SDK, then the value should be <code>Provider=MSOLAP;Data Source=local</code> . If you use v1.1, then it should be <code>Local Analysis Server</code> .
Data Source Information: Catalog	Specify the list of catalogs available, if you imported data from your data source. The cube tables correspond to the catalog you use in the connection pool.
System IP or Hostname	This property is only displayed for SAP/BW data sources. Provide the host name or IP address of the SAP data server. This field corresponds to the parameter <code>ashost</code> in the SAP/BW connect string.
System Number	This property is only displayed for SAP/BW data sources. Provide the SAP system number. This is a two-digit number assigned to an SAP instance, also called Web Application Server, or WAS. This field corresponds to the parameter <code>sysnr</code> in the SAP/BW connect string.
Client Number	This property is only displayed for SAP/BW data sources. Provide the SAP client number. This is a three-digit number assigned to the self-contained unit called Client in SAP. A Client can be a training, development, testing, or production client, or it can represent different divisions in a large company. This field corresponds to the parameter <code>client</code> in the SAP/BW connect string.
Language	This property is only displayed for SAP/BW data sources. Provide the SAP language code used when logging in to the data source (for example, EN for English or DE for German). This field corresponds to the parameter <code>lang</code> in the SAP/BW connect string.
Additional Parameters	This property is only displayed for SAP/BW data sources. Optionally, provide additional connection string parameters in the format <code>param=value</code> . Delimit multiple parameters with a colon.
Use session	An option that controls whether queries go through a common session. Consult your multidimensional data source administrator to determine whether this option should be enabled. Default is Off (not selected).

Setting Connection Pool Properties in the Connection Scripts Tab

This section describes the properties in the Connection Scripts tab of the Connection Pool dialog. The Connection Scripts tab is available for ODBC, OCI, Oracle OLAP, ADF, and DB2 data sources.

You can create connection scripts and set them to be run before the connection is established, before a query is run, after a query is run, or after the connection is disconnected. For example, you can create a connection script that, on connect, inserts the name of the user and the connection time into a table.

Connection scripts can contain any commands accepted by the database, such as a command to turn on quoted identifiers. In a mainframe environment, a script could be used to set the secondary authorization ID when connecting to DB2 to force a security

exit to a mainframe security package such as RACF. This enables mainframe environments to maintain security in one central location.

Because the connection script is sent directly to the data source, the script should use native SQL or another language understood by the data source, not Oracle BI Server Logical SQL.

To create connection scripts for data sources:

- In the Connection Pool dialog, click the Connection Scripts tab, and then complete the fields using the information in [Table 6-5](#).

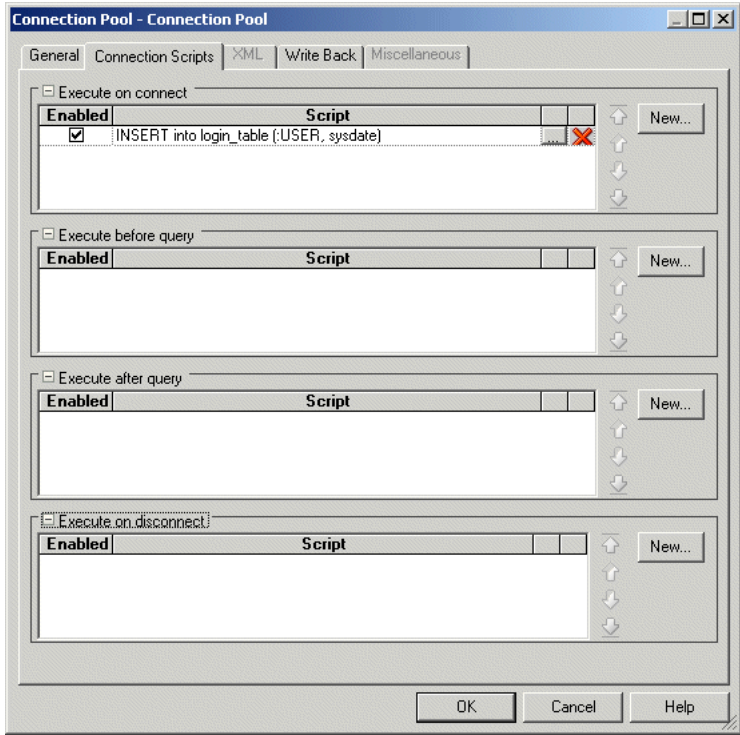
To enter a new connection script, click **New** next to the appropriate script type. Then, enter or paste the SQL statements for the script and click **OK**.

You can edit existing scripts by clicking the ellipsis button to launch the Physical SQL window. Use the Up Arrow and Down Arrow buttons to reorder existing scripts.

Click **Delete** to remove a script.

[Figure 6-3](#) shows the Connection Scripts tab of the Connection Pool dialog.

Figure 6-3 Connection Scripts Tab of the Connection Pool Dialog



[Table 6-5](#) describes the properties in the Connection Scripts tab of the Connection Pool dialog.

Table 6-5 Connection Pool Properties in the Connection Scripts Tab

Property	Description
Execute on connect	Contains SQL queries that are executed before the connection is established.
Execute before query	Contains SQL queries that are executed before the query is run.

Table 6–5 (Cont.) Connection Pool Properties in the Connection Scripts Tab

Property	Description
Execute after query	Contains SQL queries that are executed after the query is run.
Execute on disconnect	Contains SQL queries that are executed after the connection is closed.

Setting Connection Pool Properties in the XML Tab

This section describes the properties in the XML tab of the Connection Pool dialog. See [Table 6–6](#) for details. The XML tab is only available for XML and XML Server data sources.

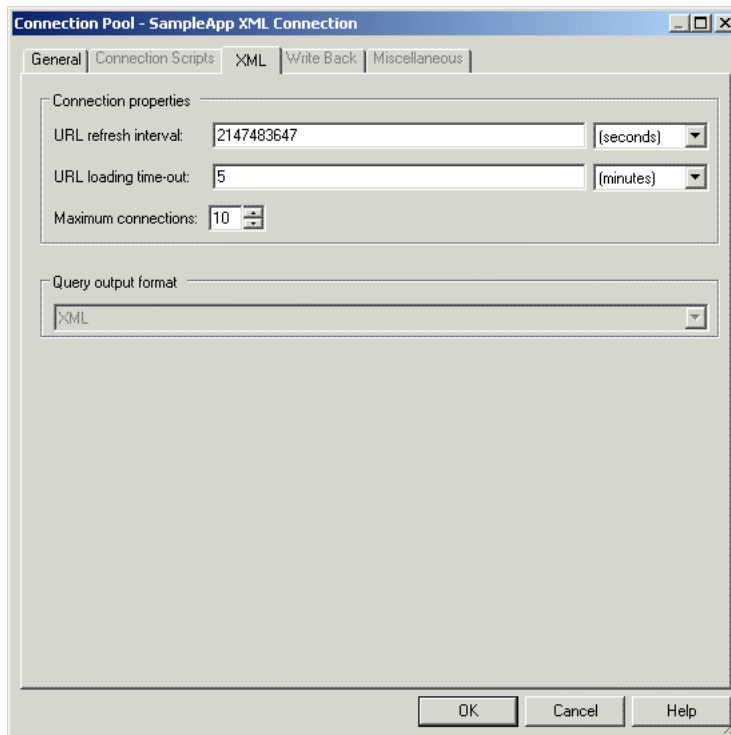
Caution: The XML tab of the Connection Pool dialog provides the same functionality as the XML tab of the Physical Table dialog. However, the properties in the XML tab of the Physical Table dialog override the corresponding settings in the Connection Pool dialog.

To set connection pool properties for XML data sources:

- In the Connection Pool dialog, click the XML tab, and then complete the fields using the information in [Table 6–6](#).

[Figure 6–4](#) shows the XML tab of the Connection Pool dialog.

Figure 6–4 XML Tab of the Connection Pool Dialog



[Table 6–6](#) describes the properties in the XML tab of the Connection Pool dialog.

Table 6–6 Connection Pool Properties in the XML Tab

Property	Description
Connection method: Search script	This property is only displayed for XML Server data sources. Click Browse to locate the appropriate search script.
Connection properties: URL refresh interval	This property is used for XML data sources and is not available for XML Server data sources. The refresh interval is analogous to setting cache persistence for database tables. The URL refresh interval is the time interval after which the XML data source is queried again directly rather than using results in cache. The default setting is infinite, meaning the XML data source is never refreshed. If you specified a URL to access the data source, set the URL refresh interval. <ul style="list-style-type: none"> ■ Select a value from the list (Infinite, Days, Hours, Minutes or Seconds). ■ Specify a whole number as the numeric portion of the interval.
Connection properties: URL loading time-out	The timeout interval for queries. The default is 15 minutes. If you specified a URL to access the data source, set the URL loading time-out as follows: <ul style="list-style-type: none"> ■ Select a value from the list (Infinite, Days, Hours, Minutes or Seconds). ■ Specify a whole number as the numeric portion of the interval.
Connection properties: Maximum connections	The maximum number of connections. The default is 10.
Query input supplements: Header file/Trailer file	This property is only displayed for XML Server data sources. Click Browse to locate the header and trailer files.
Query output format	For XML data sources, choose only XML . Other output formats are available for XML Server data sources.

Setting Connection Pool Properties in the Write Back Tab

This section describes the properties in the Write Back tab of the Connection Pool dialog. The Write Back tab is available for ODBC, OCI, Oracle OLAP, ADF Business Component, and DB2 data sources.

To set write-back properties for data sources:

- In the Connection Pool dialog, click the Write Back tab, and then complete the fields using the information in [Table 6–7](#).

[Figure 6–3](#) shows the Write Back tab of the Connection Pool dialog.

Figure 6–5 Write Back Tab of the Connection Pool Dialog

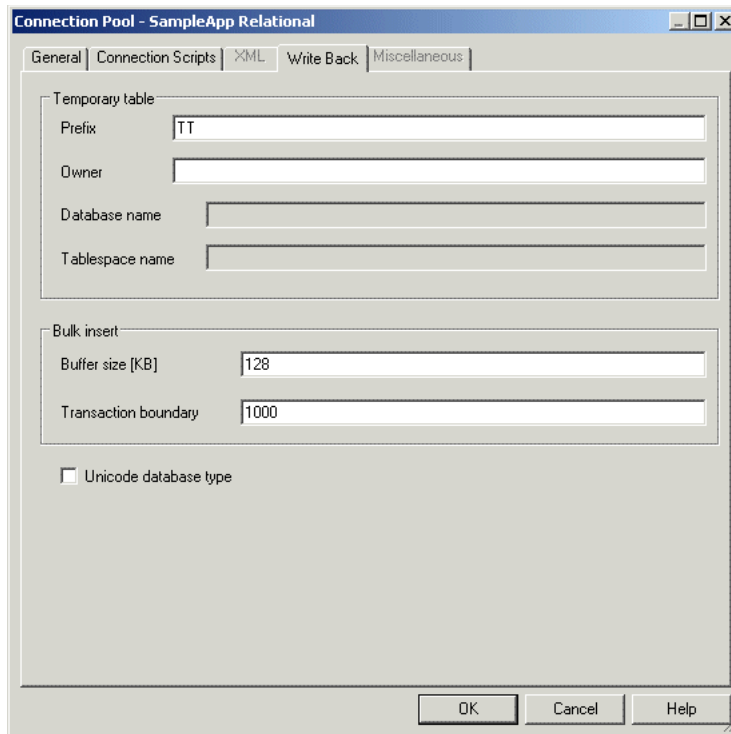


Table 6–7 describes the properties in the Write Back tab of the Connection Pool dialog.

Table 6–7 Connection Pool Properties in the Write Back Tab

Property	Description
Temporary table: Prefix	When the Oracle BI Server creates a temporary table, these are the first two characters in the temporary table name. The default value is TT.
Temporary table: Owner	Table owner name used to qualify a temporary table name in a SQL statement, for example to create the table <code>owner.tablename</code> . If left blank, the user name specified in the writeable connection pool is used to qualify the table name and the Shared logon field on the General tab should also be set.
Temporary table: Database name	Database where the temporary table will be created. This property applies only to IBM OS/390 because IBM OS/390 requires database name qualifier to be part of the <code>CREATE TABLE</code> statement. If left blank, OS/390 defaults the target database to a system database for which the users may not have Create Table privileges.
Temporary table: Tablespace name	Tablespace where the temporary table will be created. This property applies to OS/390 only as OS/390 requires tablespace name qualifier to be part of the <code>CREATE TABLE</code> statement. If left blank, OS/390 defaults the target database to a system database for which the users may not have Create Table privileges.
Bulk insert: Buffer size (KB)	Used for limiting the number of bytes each time data is inserted in a database table. For optimum performance, consider setting this parameter to 128. See "About Setting the Buffer Size and Transaction Boundary" for additional information.

Table 6–7 (Cont.) Connection Pool Properties in the Write Back Tab

Property	Description
Bulk insert: Transaction boundary	Controls the batch size for an insert in a database table. For optimum performance, consider setting this parameter to 1000. See " About Setting the Buffer Size and Transaction Boundary " for additional information.
Unicode database type	<p>Select this option when working with columns of an explicit Unicode data type, such as NCHAR, in a Unicode database. This makes sure that the binding is correct and that data is inserted correctly. Different database vendors provide different character data types and different levels of Unicode support. Use the following general guidelines to determine when to set this option:</p> <ul style="list-style-type: none"> ■ On a database where CHAR data type supports Unicode and there is no separate NCHAR data type, do not select this option. ■ On a database where NCHAR data type is available, it is recommended to select this option. ■ On a database where CHAR and NCHAR data type are configured to support Unicode, selecting this option is optional. <p>Note: Unicode and non-Unicode data types cannot coexist in a single non-Unicode database. For example, mixing the CHAR and NCHAR data types in a single non-Unicode database environment is not supported.</p>

Setting Connection Pool Properties in the Miscellaneous Tab

This section describes the properties in the Miscellaneous tab of the Connection Pool dialog. The Miscellaneous tab is only available for ADF Business Component data sources.

To set application properties for ADF Business Component data sources:

- In the Connection Pool dialog, click the Miscellaneous tab, and then complete the fields using the information in [Table 6–8](#).

[Figure 6–6](#) shows the Miscellaneous tab of the Connection Pool dialog.

Figure 6–6 Miscellaneous Tab of the Connection Pool Dialog




Table 6–8 describes the properties in the Miscellaneous tab of the Connection Pool dialog.

Table 6–8 Connection Pool Properties in the Miscellaneous Tab

Property	Description
AppModule Definition	The fully qualified Java package name of the Root Application Module to which you want to connect, such as <code>oracle.apps.fii.receivables.model.RootAppModule</code> .
AppModule Config	Determines which application configuration is used in the connection, such as <code>RootAppModuleShared</code> .
URL	The URL to the Oracle Business Intelligence broker servlet, in the format: <code>http://host:port/APP_DEPLOYMENT_NAME/obieebroker</code> For example: <code>http://localhost:7001/SnowflakeSalesApp/obieebroker</code> The URL is case-sensitive.

Table 6–8 (Cont.) Connection Pool Properties in the Miscellaneous Tab

Property	Description
SQL Bypass Database	<p>(Optional) The name of the SQL Bypass database. The SQL Bypass database must be a physical database in the Physical layer of the repository. The database object for the SQL Bypass database must have a valid connection pool, with connection information that points to the same database that is being used by the JDBC Data source defined in the WebLogic Server.</p> <p>The SQL Bypass database does not need to have any tables under it. After a valid database name is supplied, the SQL Bypass feature is enabled for all queries.</p> <p>The SQL Bypass feature directly queries the database so that aggregations and other transformations are pushed down where possible, reducing the amount of data streamed and worked on in Oracle Business Intelligence. See "About Specifying a SQL Bypass Database" for more information.</p>

Setting Up Persist Connection Pools

A persist connection pool is a database property that is used for specific types of queries (typically used to support Marketing queries). In some queries, all of the logical query cannot be sent to the transactional database because that database might not support all of the functions in the query. This issue might be solved by temporarily constructing a physical table in the database and rewriting the Oracle BI Server query to reference the new temporary physical table.

You can use the persist connection pool in the following situations:

- **Populate stored procedures.** Use to rewrite the Logical SQL result set to a managed table. Typically used by Oracle's Siebel Marketing Server to write segmentation cache result sets.
- **Perform a generalized subquery.** Stores a nonfunction subquery in a temporary table, and then rewrites the original subquery result against this table. Reduces data movement between the Oracle BI Server and the database, supports unlimited IN list values, and might result in improved performance.

In these situations, the user issuing the Logical SQL query must have been granted the Populate privilege on the target database.

The persist connection pool functionality designates a connection pool with write-back capabilities for processing this type of query. You can assign one connection pool in a single database as a persist connection pool. If this functionality is enabled, the user name specified in the connection pool must have the privileges to create DDL (Data Definition Language) and DML (Data Manipulation Language) in the database.

To assign a persist connection pool:

1. In the Physical layer of the Administration Tool, double-click the database object for which you want to assign a persist connection pool.
2. In the Database dialog, click the General tab.
3. In the **Persist connection pool** area, click **Set**.

If there is only one connection pool, it appears in the **Persist connection pool** field.

4. If there are multiple connection pools, in the Browse dialog, select the appropriate connection pool, and then click **OK**.

The selected connection pool name appears in the **Persist connection pool** field.

5. (Optional) To set write-back properties, click the Connection Pools tab.
6. In the connection pool list, double-click the connection pool.
7. In the Connection Pool dialog, click the Write Back tab.
8. Complete the fields using [Table 6-7](#) as a guide. See also "[About Setting the Buffer Size and Transaction Boundary](#)" for additional information.
9. Click **OK**, then click **OK** again to save the persist connection pool.

To remove a persist connection pool:

1. In the Physical layer of the Administration Tool, double-click the database object that contains the persist connection pool you want to remove.
2. In the Database dialog, click the General tab.
3. In the **Persist connection pool** area, click **Clear**.
The database name is replaced by `not assigned` in the **Persist connection pool** field.
4. Click **OK**.

About Setting the Buffer Size and Transaction Boundary

If each row size in a result set is 1 KB and the buffer size is 20 KB, then the maximum array size is 20 KB. If there are 120 rows, there are 6 batches with each batch size limited to 20 rows.

If you set **Transaction boundary** to 3, the server commits twice. The first time, the server commits after row 60 (3 * 20). The second time, the server commits after row 120. If there is a failure when the server commits, the server only rolls back the current transaction. For example, if there are two commits and the first commit succeeds but the second commit fails, the server only rolls back the second commit.

For optimum performance, consider setting the buffer size to 128 and the transaction boundary to 1000.

Working with Physical Tables, Cubes, and Joins

The Physical layer of the Oracle BI repository contains objects that represent physical data constructs from back-end data sources. The Physical layer defines the objects and relationships available to the Oracle BI Server for writing physical queries. This layer encapsulates data source dependencies to enable portability and federation.

Each data source of the repository model typically has its own discrete physical model in the Physical layer. The top-level object in the Physical layer is a "database," and the type of database determines which features and rules apply to that physical model. For example, a relational database such as "Oracle 11g" has relational objects such as physical tables and joins. In contrast, a multidimensional source such as "Essbase 9" has cube tables and physical hierarchies. Therefore, some sections of this chapter apply to only certain database types.

Physical tables, cubes, joins, and other objects in the Physical layer are typically created automatically when you import metadata from your data sources. After these objects have been imported, you can perform tasks like creating additional join paths that are not in the data source, create alias tables for physical tables that need to serve in different roles, and adjust properties of physical hierarchies from multidimensional data sources.

This chapter contains the following topics:

- [Working with the Physical Diagram](#)
- [Creating Physical Layer Folders](#)
- [Working with Physical Tables](#)
- [Working with Multidimensional Sources in the Physical Layer](#)
- [Working with Essbase Data Sources](#)
- [Working with Hyperion Financial Management Data Sources](#)
- [Working with Oracle OLAP Data Sources](#)
- [Working with Physical Foreign Keys and Joins](#)
- [Deploying Opaque Views](#)
- [Using Hints](#)
- [Displaying and Updating Row Counts for Physical Tables and Columns](#)

Working with the Physical Diagram

In addition to working with Physical layer objects in the right pane of the Administration Tool, you can open the Physical Diagram view to see a graphical model of tables and joins.

Note: The Physical Diagram is typically used with relational and XML sources rather than multidimensional sources. Although the Physical Diagram view for a multidimensional source does display a denormalized table representation of a cube table, the primary means of working with a multidimensional physical model is by working in the physical tree using dimensions, hierarchies and columns.


To access the Physical Diagram, right-click an object in the Physical layer tree view (such as a physical database or table) and select **Physical Diagram**. Then, select one of the following options:

- **Selected Object(s) Only.** Displays only the selected objects. Joins appear only if they exist between the objects that you select.
- **Object(s) and Direct Joins.** Displays the selected objects and any tables that join to the objects that you select.
- **Object(s) and All Joins.** Displays the selected objects, as well as each object that is related directly or indirectly to the selected object through some join path. If all the objects in a schema are related, then using this option diagrams every table, even if you only select one table.

Note that the Physical Diagram displays only physical tables and joins. It does not display other Physical layer objects, such as connection pools, physical hierarchies, or levels.

[Figure 7–1](#) shows the Physical Diagram.

Figure 7–1 Physical Diagram



You can also open the Physical Diagram by selecting one or more objects in the tree view and then clicking the **Physical Diagram** button on the toolbar:



Only the objects you selected appear. Joins appear only if they exist between the selected objects. Joins are represented by a line with an arrow at the "one" end of the join.

To help you better understand the logical-to-physical mappings in your model, you can view the physical objects that are associated with a particular logical object by selecting one or more business models, logical tables, or logical table sources in the Business Model and Mapping layer tree view and then clicking the **Physical Diagram** button on the toolbar. Only physical objects that are related to the objects you selected appear. You can view the same information by right-clicking a logical object and selecting **Objects and Direct Join(s) within Business Model** from the Physical Diagram submenu. You can also choose one of the other Physical Diagram display options.

To add additional tables to the Physical Diagram, leave the Physical Diagram window open and then right-click the table or tables you want to add. Then, select **Physical Diagram** and choose one of the display options.

Additional options are available in the right-click menu for the graphical tables and joins displayed in the Physical Diagram. For example, you can delete objects or view their properties, or you can add additional related objects using the right-click options **Add Direct Joins**, **Add Tables Joined to Whole Selection**, and **Add All Joins**. You can also select **Find in Tree View** to locate a particular object in the Physical layer tree view in the right pane, or check out objects in online mode.

You can also right-click an object in the Physical Diagram view and select **Hide** to hide particular objects in the diagram. Note that this effect is temporary and does not persist.

See also the following sections:

- ["Using the Physical and Business Model Diagrams"](#) for information about zooming, panning, and controlling the layout of the tables
- ["Defining Physical Joins with the Physical Diagram"](#) for information about defining physical joins

Creating Physical Layer Folders

This section contains the following topics:

- [Creating Physical Layer Catalogs and Schemas](#)
- [Using a Variable to Specify the Name of a Catalog or Schema](#)
- [Setting Up Display Folders in the Physical Layer](#)

Creating Physical Layer Catalogs and Schemas

Catalogs are optional ways to group different schemas. A catalog contains all the schemas (metadata) for a physical database object. A schema contains only the metadata information for a particular user or application. Model the Physical layer after the way your data source is structured.

Note the following:

- You must create a physical database object before you can create a physical catalog object or a physical schema object.
- After you implement a certain type of grouping, you cannot change it later. For example, if you decide to implement database > schema > table, you cannot add a catalog afterward.

Creating Catalogs

In the Physical layer of a large repository, administrators can create physical catalogs that contain one or more physical schemas.

To create a catalog:

1. In the Physical layer of the Administration Tool, right-click a physical database and select **New Object**, then select **Physical Catalog**.
2. In the Physical Catalog dialog, type a name for the catalog.
3. Type a description for the catalog, and then click **OK**.

Creating Schemas

The schema object contains tables and columns for a physical schema. Schema objects are optional in the Physical layer of the Administration Tool.

To create a schema:

1. In the Physical layer of the Administration Tool, right-click a physical database or physical catalog and select **New Object**, then select **Physical Schema**.
2. In the Physical Schema dialog, type a name.
3. Type a description for the schema, and then click **OK**.

Using a Variable to Specify the Name of a Catalog or Schema

You can use a variable to specify the names of catalog and schema objects. For example, you have data for multiple clients and you structured the data source so that data for each client was in a separate catalog. You would initialize a session variable named `Client`, for example, that could be used to set the name for the catalog object dynamically when a user signs on to the Oracle BI Server.

You specify the session variable to use in the Dynamic Name tab of the Physical Catalog or Physical Schema dialog.

Note: The Dynamic Name tab is not active unless at least one session variable is defined.

To specify the session variable to use in the Dynamic Name tab:

1. In the **Name** column of the Dynamic Name tab, click the name of the session variable that you want to use. The initial value for the variable (if any) is shown in the **Default Initializer** column.
2. To select the highlighted variable, click **Select**.

The name of the variable is displayed in the dynamic name field, and the **Select** button toggles to the **Clear** button.

To remove assignment for a session variable in the Dynamic Name tab:

- Click **Clear** to remove the assignment for the variable as the dynamic name.

The value **not assigned** is displayed in the dynamic name field, and the **Clear** button toggles to the **Select** button.

To sort column entries in the Dynamic Name tab:

- You can sort the entries in a column by clicking the **Name** or **Default Initializer** column heading. Clicking a column heading toggles the order of the entries in that column between ascending and descending order, according to the column type.

Setting Up Display Folders in the Physical Layer

You can create display folders to organize table objects in the Physical layer. They have no effect on query processing. After you create a display folder, the selected tables appear in the folder as a shortcut and in the Physical layer tree as an object. You can hide the objects so that you only view the shortcuts in the display folder. See the information about the Repository tab of the Options dialog in "[Setting Preferences](#)" for more information about hiding these objects.

Note: Deleting a table in a display folder deletes only the shortcut to that object. When you delete a column in a display folder, however, the column is actually deleted.

To set up a physical display folder:

1. In the Physical layer of the Administration Tool, right-click a physical database and select **New Object**, then select **Physical Display Folder**.
2. In the Physical Display Folder dialog, type a name for the folder.
3. To add tables to the display folder, click **Add**. Then, in the Browse dialog, select the fact or physical tables you want to add to the folder and click **Select**.
Alternatively, you can drag one or more physical tables to the display folder after you close the dialog.
4. Click **OK**.

Working with Physical Tables

This section explains how to work with physical table objects in the Physical layer of the Oracle BI repository.

Note that both physical tables from relational data sources and physical cube tables from multidimensional data sources both use the Physical Table table type. Many of the tasks described in this section apply to both relational and multidimensional data sources. See also "[Working with Multidimensional Sources in the Physical Layer](#)" for additional information specific to multidimensional data sources.

This section contains the following topics:

- [About Tables in the Physical Layer](#)
- [About Physical Alias Tables](#)
- [Creating and Managing Physical Tables and Physical Cube Tables](#)
- [Creating and Managing Columns and Keys for Relational and Cube Tables](#)
- [Viewing Data in Physical Tables or Columns](#)

About Tables in the Physical Layer

A physical table is an object in the Physical layer of the Oracle BI repository that corresponds to a table in a data source. Metadata for physical tables is usually imported from the data source. This metadata enables the Oracle BI Server to access the data source tables with SQL requests.

When you delete a physical table, all dependent objects are deleted (for example, columns, keys, and foreign keys). When you delete a physical cube table, hierarchies are also deleted. The deletion fails if an alias exists on the physical table.

In addition to importing data source tables into the Physical layer, you can create virtual physical tables in the Physical layer, using values in the **Table Type** field in the Physical Table dialog. Creating virtual tables can provide the Oracle BI Server and the underlying data sources with the proper metadata to perform some advanced query requests.

A virtual physical table can be a stored procedure, or a `SELECT` statement. A virtual physical table created from a `SELECT` statement is also called an opaque view. You can

define an opaque view, and then deploy it in your data source to create a deployed view. See ["Deploying Opaque Views"](#) for more information.

Use the **Table Type** list in the General tab of the Physical Table dialog to specify the physical table object type. [Table 7-1](#) describes the available object types.

Table 7-1 Table Types for Physical Tables

Table Type	Description
Physical Table	Specifies that the physical table object represents a data source table. This table type is used for both relational physical tables and multidimensional cube tables.
Stored Proc	<p>Specifies that the physical table object is a stored procedure. When you select this option, you type the stored procedure in the text box. Requests for this table will call the stored procedure.</p> <p>For stored procedures that are data source-specific, select Use database specific SQL. When you select this option, the Database column displays supported data sources by brand, with Default as the root. You can enter data source-specific initialization strings by selecting the database type on the left and entering the corresponding string on the right. The initialization string for the Default option is run when the queried database type does not have a corresponding database-specific string defined.</p> <p>Stored procedures within an Oracle Database do not typically return result sets. Therefore, they cannot be initiated from within Oracle Business Intelligence. You need to rewrite the procedure as an Oracle function, use it in a <code>SELECT</code> statement in the Administration Tool initialization block, and then associate it with the appropriate Oracle BI Server session variables in the Session Variables dialog.</p> <p>The following example shows a SQL initialization string using the <code>GET_ROLES</code> function that is associated with the <code>USER</code>, <code>GROUP</code>, and <code>DISPLAYNAME</code> variables. The function takes a user Id as a parameter and returns a semicolon-delimited list of group names:</p> <pre>SELECT user_id, get_roles(user_id), first_name ' ' last_name FROM csx_security_table WHERE user_id = ':USER' and password = ':PASSWORD'</pre>
Select	<p>Specifies that the physical table object is a <code>SELECT</code> statement. When you select this option, you type the <code>SELECT</code> statement in the text field, and you also need to manually create the table columns. The column names must match the ones specified in the <code>SELECT</code> statement. Column aliases are required for advanced SQL functions, such as aggregates and <code>CASE</code> statements.</p> <p>Requests for this table will execute the <code>SELECT</code> statement.</p> <p>For <code>SELECT</code> statements that are data source-specific, select Use database specific SQL. When you select this option, the Database column displays supported data sources by brand, with Default as the root. You can enter data source-specific initialization strings by selecting the database type on the left and entering the corresponding string on the right. The initialization string for the Default option is run when the queried database type does not have a corresponding database-specific string defined.</p> <p>This type of table is also called an opaque view. See "Deploying Opaque Views" for more information.</p>

About Physical Alias Tables

An alias table (alias) is a physical table that references a different physical table as its source (called the original table). Alias tables can be an important part of designing a Physical layer because they enable you to reuse an existing table more than once, without having to import it several times.

There are two main reasons to create an alias table:

- To set up multiple tables, each with different keys, names, or joins, when a single data source table needs to serve in different semantic roles. Setting up alias tables in this case is a way to avoid triangular or circular joins.

For example, an order date and a shipping date in a fact table may both point to the same column in the time dimension data source table, but you should alias the dimension table so that each role is presented as a separately labeled alias table with a single join. These separate roles carry over into the business model, so that "Order Date" and "Ship Date" are part of two different logical dimensions. If a single logical query contains both columns, the physical query uses aliases in the SQL statement so that it can include both of them.

You can also use aliases to enable a data source table to play the role of both a fact table, and a dimension table that joins to another fact table (often called a "fan trap").

- To include best practice naming conventions for physical table names. For example, you can prefix the alias table name with the table type (such as fact, dimension, or bridge), and leave the original physical table names as-is. Some organizations alias all physical tables to enforce best practice naming conventions. In this case, all mappings and joins are based on the alias tables rather than the original tables.

Alias table names appear in physical SQL queries. Using alias tables to provide meaningful table names can make SQL queries referencing those tables easier to read. For example:

```
WITH
SAWITH0 AS (select sum(T835.Dollars) as c1
from
    FactsRevT835/*AllRevenue(Billed Time Join)*/)
select distinct 0 as c1,
    D1.c1 as c2
from
    SAWITH0 D1
order by c1
```

In this query, the meaningful alias table name "A11 Revenue (Billed Time Join)" has been applied to the terse original physical table name "FACTSREV." In this case, the alias table name provides information about which role the table was playing each time it appears in SQL queries.

Alias tables can have cache properties that differ from their original tables. To set different cache properties for an alias table, select the option **Override Source Table Caching Properties** in the Physical Table dialog for the alias table. In alias tables, columns cannot be added, deleted, or modified. Because columns are automatically synchronized, no manual intervention is required.

Synchronization ensures that the original tables and their related alias tables have the same column definitions. For example, if you delete a column in the original table, the column is automatically removed from the alias table.

You cannot delete an original table unless you delete all its alias tables first. Alternatively, you can select the original table and all its alias tables and delete them at the same time.

You can change the original table of an alias table, if the new original table is a superset of the current original table. However, this could result in an inconsistent repository if changing the original table deletes columns that are being used. If you attempt to do this, a warning message appears to let you know that this could cause a problem and lets you cancel the action. Running a consistency check identifies orphaned aliases.

When you edit a physical table or column in online mode, all alias tables and columns must be checked out. The behavior of online checkout uses the following conventions:

- If an original table or column is checked out, all its alias tables and columns are checked out.
- If an alias table or column is checked out, its original table and column are checked out.
- The checkout option is available for online repositories (if not read-only) and for all original and alias tables and columns.

Alias tables inherit some properties from their original tables. A property that is proxied is a value that is always the same as the original table, and cannot be changed. (In other words, the proxied properties are the ones that are dimmed in the alias table dialog.) If the original table changes its value for that particular property, the same change is applied on the alias table.

The following is a list of the properties that are proxied:

- Cacheable (the inherited property can be overridden)
- Cache never expires and Cache persistence time (the inherited properties can be overridden)
- Row Count
- Last Updated
- Table Type
- External Db Specifications

The following is list of the properties that are not proxied:

- Name
- Description
- Display Folder Containers
- Foreign Keys
- Columns

Note: Alias tables and original tables never share columns. Aliases and original tables have distinctly different columns that alias each other.

- Table Keys
- Complex Joins

- Source Connection Pool
- Polling Frequency
- All XML attributes

Creating and Managing Physical Tables and Physical Cube Tables

Use the General tab of the Physical Table dialog to create or edit physical tables and physical cube tables in the Physical layer of the Administration Tool.

This section contains the following topics:

- [Creating or Editing Physical Tables](#)
- [Creating Alias Tables](#)
- [Setting Physical Table Properties for XML Data Sources](#)

Creating or Editing Physical Tables

This section describes how to create or edit the general properties for a table, including both relational physical tables and physical cube tables.

To create a physical table or edit general properties for tables:

1. In the Physical layer of the Administration Tool, perform one of the following steps:
 - To create a physical table, right-click the physical database or physical catalog and select **New Object**, then select **Physical Table**.
If your database object has physical schemas defined, right-click the physical schema and select **New Physical Table**.
 - To create a physical cube table for a multidimensional data source, right-click the physical database and select **New Object**, then select **Cube Table**.
-
- Caution:** It is strongly recommended that you import cube tables, not create them manually.
-
- To edit an existing physical table, double-click the physical table object in the Physical layer.
2. In the Physical Table dialog, complete the fields using [Table 7-2](#) as a guide.

Table 7-2 *General Properties for Physical Tables*

Property	Description
Name	The name of the physical table.
Table Type	Physical Table values: Physical Table, Stored Proc (stored procedure), or Select. Physical Cube Table values: Physical Table or Select. See Table 7-1 for more information.

Table 7–2 (Cont.) General Properties for Physical Tables

Property	Description
Use Dynamic Name	<p>Select this option to use a session variable to specify the physical table name, similar to catalog and schema objects. This option is available for non-multidimensional data source tables when you select a table type of Physical Table.</p> <p>You might want to choose this option if you have a multi-tenancy implementation and you want to define a separate physical table name for each customer. Another example would be to select between primary and shadow tables that are valid at different times in your ETL cycle. In both cases, you can assign session variables to dynamically select the appropriate table.</p>
Default Initialization String / Use database specific SQL	<p>For non-multidimensional data source tables (not alias tables), this option appears if you choose a Table Type of Stored Proc or Select. For multidimensional data source tables, this appears if you choose a Table Type of Select.</p> <p>When you select this option, you can specify the data source and type the SQL statements.</p> <p>See Table 7–1 for more information.</p>
Cacheable	<p>Select this option to include the table in the Oracle BI Server query cache. Typically, you should select this option for tables that do not need to be accessed in real time.</p> <p>When you select this option, the Cache persistence time settings become active.</p> <p>Note that there are additional configuration settings that affect the behavior of the query cache. See "Configuring Query Caching" in <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i> for full information.</p>
Cache never expires	<p>When you select this option, cache entries do not automatically expire. This could be useful when a table is important to a large number of queries users might run. For example, if most of your queries have a reference to an account object, keeping it cached indefinitely could actually improve performance rather than compromise it.</p> <p>Note that selecting this option does not mean that an entry always remains in the cache. Other invalidation techniques, such as manual purging, LRU (Least Recently Used) replacement, metadata changes, or use of the cache polling table can result in entries being removed from the cache.</p>
Cache persistence time	<p>How long table entries should persist in the query cache, or in other words, the cache expiration time.</p> <p>Setting a cache persistence time is useful for OLTP data sources and other data sources that are updated frequently. For example, you could set this option refresh the underlying physical tables daily for a particular dashboard.</p> <p>If a query references multiple physical tables with different persistence times, the cache entry for the query exists for the shortest persistence time set for any of the tables referenced in the query. This makes sure that no subsequent query gets a cache hit from an expired cache entry.</p> <p>For more information, see "Troubleshooting Problems with Event Polling Tables" in <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>.</p>

Table 7–2 (Cont.) General Properties for Physical Tables

Property	Description
External name	Applies to physical cube tables from multidimensional data sources. The external name is the physical name that is used when referencing the cube table in physical SQL queries. This value must reflect the external name defined in the data source.
Display Column	For Essbase data sources only. See "Working with Essbase Data Sources" for more information.
Hint	Available only for some data sources. See "Using Hints" for more information.

Creating Alias Tables

To create an alias table, right-click an existing physical table and select **New Object**, then select **Alias**. You can also create aliases on opaque views and stored procedures.

[Table 7–3](#) describes properties that are specific to alias tables. Refer to [Table 7–2](#) for information about other table properties that are shared between physical tables and alias tables.

Table 7–3 Properties Specific to Physical Alias Tables

Property	Description
Source Table	Applies to alias tables. Click Select to choose the original physical table from which to create an alias table.
Override Source Table Caching Properties	Option available for alias tables. When selected, the cacheable properties become available and you can clear or select the appropriate options.

Setting Physical Table Properties for XML Data Sources

Use the XML tab to set or edit properties for an XML data source. The XML tab of the Physical Table dialog provides the same functionality as the XML tab of the Connection Pool dialog. However, setting properties in the Physical Table dialog overrides the corresponding settings in the Connection Pool dialog. See ["Setting Connection Pool Properties in the XML Tab"](#) for more information.

About the Calc Scripts Tab for Essbase Data Sources

The Calc Scripts tab for Essbase data sources is reserved for a future release.

Creating and Managing Columns and Keys for Relational and Cube Tables

Each physical table and physical cube table in the Physical layer of the Administration Tool has one or more physical columns. You can use the Columns, Keys, and Foreign Keys tabs in the Physical Table dialog to view, create new, and edit existing columns, keys, and foreign keys that are associated with the table.

The following list describes the buttons that appear in the tabs:

- **New.** Lets you create a new object by opening the dialog that corresponds to the tab.
- **Edit.** When you select an object and then click **Edit**, the dialog that corresponds to the tab appears. You can then edit the properties of the object.
- **Delete.** Deletes the selected object.

This section contains the following topics:

- [Creating and Editing a Column in a Physical Table](#)
- [Specifying a Primary Key for a Physical Table](#)
- [Deleting Physical Columns for All Data Sources](#)

Creating and Editing a Column in a Physical Table

If the column is imported, the properties of the column are set automatically. The following list contains information about nullable and data type values for columns imported into the Physical layer.

- **Nullable.** Indicates whether null values are allowed for the column. If null values can exist in the underlying table, you need to select this option. This allows null values to be returned to the user, which is expected with certain functions and with outer joins. It is generally safe to change a non-nullable value to a nullable value in a physical column.
- **Type.** Indicates the data type of the column. Use caution when changing the data type. Setting the values to ones that are incorrect in the underlying data source might cause unexpected results. If there are any data type mismatches, correct them in the repository or reimport the columns that have mismatched data types.

If you reimport columns, you also need to remap any logical column sources that reference the remapped columns. The data type of a logical column in the business model must match the data type of its physical column source. The Oracle BI Server passes these logical column data types to client applications.

Longvarchar and longvarbinary data types are supported for writing complete Logical SQL statements into usage tracking tables for debugging purposes. They are not supported for general-purpose queries, and cannot be displayed in Presentation Services. Use direct SQL utilities to access columns with these data types.

Except when stated otherwise, the characteristics and behavior of a physical cube column are the same as for other physical columns.

Note: Creating, modifying, or deleting a column in an original physical table also creates, modifies, or deletes the same column on all its alias tables.

To create or edit a physical column:

1. In the Physical layer of the Administration Tool, perform one of the following steps:
 - To create a physical column, right-click a physical table and select **New Object**, then select **Physical Column**.
 - To create a physical cube column for a multidimensional data source, right-click a physical cube table and select **New Object**, then select **Physical Cube Column**.
 - To edit an existing physical column, double-click the physical column object in the Physical layer.
2. In the Physical Column dialog, type a name for the physical column.

For XML data sources, this field stores and displays the unqualified name of a column (attribute) in an XML document.

3. In the **Type** field, select a data type for the physical column.
4. If applicable, specify the length of the data type.
For multidimensional data sources, if you select `VARCHAR`, you need to type a value in the **Length** field.
5. Select the **Nullable** option if the column is allowed to have null values.
6. In the **External Name** field, type an external name.
 - Required if the same name (such as `STATE`) is used in multiple hierarchies.
 - Optional for XML documents. The **External Name** field stores and displays the fully qualified name of a column (attribute).
7. (Multidimensional data sources) When the physical cube column is a measure, in the **Aggregation role** list, select the appropriate value.
A new physical cube column is created as a measure by default. See ["Working with Multidimensional Sources in the Physical Layer"](#) for information about changing this behavior.
8. Click **OK**.

Specifying a Primary Key for a Physical Table

Use the Physical Key dialog to specify the column or columns that define the primary key of the physical table.

To specify a primary key for a physical table:

1. In the Physical layer of the Administration Tool, right-click a physical table and select **Properties**.
2. In the Physical Table dialog, click the Keys tab.
3. In the Keys tab, click **New**.
4. In the Physical Key dialog, type a name for the key.
5. Select the column that defines the primary key of the physical table.
6. (Optional) Type a description for the key.
7. Click **OK**.

Deleting Physical Columns for All Data Sources

When you delete a physical column, the following occurs:

- **Multidimensional data sources.** If you delete property or key columns from a level, the association is deleted and the column changes to a measure under the parent cube table.
- **Alias tables.** Deleting a column in an original physical table deletes the same column on all its alias tables.

Viewing Data in Physical Tables or Columns

You can view the data in a physical table or an individual physical column by right-clicking the object and choosing **View Data**. In online editing mode, you must check in changes before you can use this option.

View Data is not available for physical cube tables or columns. See ["Viewing Members in Physical Cube Tables"](#) for more information.

Because the View Data feature issues a row count, it is not available for data sources that do not support row counts. See "[Displaying and Updating Row Counts for Physical Tables and Columns](#)" for more information.

Caution: View Data does not work in online mode if you set the user name and password for connection pools to :USER and :PASSWORD. In offline mode, the Set values for variables dialog appears so that you can populate :USER and :PASSWORD as part of the viewing process.

Working with Multidimensional Sources in the Physical Layer

This section provides information about physical cube tables, dimensions, and hierarchies from multidimensional data sources.

- [About Physical Cube Tables](#)
- [About Measures in Multidimensional Data Sources](#)
- [Working with Physical Dimensions and Physical Hierarchies](#)
- [Working with Cube Variables for SAP/BW Data Sources](#)
- [Viewing Members in Physical Cube Tables](#)

About Physical Cube Tables

Each cube from a multidimensional data source is set up as a physical cube table, a type of physical table. It has all the capabilities of a table, such as physical cube columns and keys (optional) and foreign keys (optional). It also has cube-specific metadata such as hierarchies and levels.

When you import the physical schema, the Oracle BI Server imports the metadata for the cube, including its metrics, hierarchies, and levels. Expanding the hierarchy object in the Physical layer reveals the levels in the hierarchy. In the Physical Cube Table dialog, the Hierarchies tab lists the dimensional hierarchies in the cube.

Each multidimensional catalog in the data source can contain multiple physical cubes. You can import the metadata for one or more of these cubes into your Oracle BI repository. Although it is possible to create a cube table manually, it is recommended that you import metadata for cube tables and their components.

If you do create cubes manually, be sure to build each cube one hierarchy at a time and test each one before building another. For example, create the time hierarchy and a measure, and then test it. When it is correct, create the geography hierarchy and test it. This helps ensure that you have set up each cube correctly, and makes it easier to identify any setup errors.

About Measures in Multidimensional Data Sources

You need to select the aggregation rule for a physical cube column carefully to make sure your measures are correct. Setting it correctly might improve performance.

Always verify aggregation rules after importing cube metadata. Typically, aggregation rules are assigned correctly when you import cube metadata. However, if a measure is a calculated measure, the aggregation rule is reported as None. Therefore, you should examine the aggregation rule for all measures after importing a cube to verify that the aggregation rule has been assigned correctly.

For all measures assigned an aggregation rule value of None, contact the multidimensional data source administrator to verify that the value of the aggregation rule is accurate. If you need to change the aggregation rule, you can change it in the Physical Cube Column dialog.

Use the following guidelines to assign the correct aggregation rule:

- If the generated physical queries to the database should send an aggregation function, such as `SUM(revenue)`, then set that function as the aggregation rule. With this setting, the Oracle BI Server typically pushes the aggregation to the database in the query, but might also perform aggregations itself in certain situations.
- If the data for this measure should not be aggregated in the query or by the Oracle BI Server, use the External Aggregation rule. It is important to choose this setting when the measure uses a more complex calculation inside the data source than the Oracle BI Server can replicate with a simple aggregation rule (such as calculations for ratios, consolidations and allocations). This option is also useful when the cube persists a full set of pre-aggregated results.

About Externally Aggregated Measures

In a multidimensional data source, some cubes contain very complex, multi-level based measures. If you assign an aggregation rule of External Aggregation, the Oracle BI Server bypasses its internal aggregation mechanisms and uses the pre-aggregated measures. When imported, these measures are assigned an aggregate value of None.

The following are some guidelines for working with pre-aggregated measures:

- External aggregation only applies to multidimensional data sources (such as Essbase, Hyperion Financial Management, MS Analysis Services, and SAP/BW) that support these complex calculations.
- You cannot assign external aggregation to measures from non-multidimensional data sources. If the required aggregation rule is supported by the Oracle BI Server and can be mapped to a relational data source, then it is not complex and does not require external aggregation.
- There is only one aggregation rule for a logical measure. Therefore, a single logical column cannot federate a noncomplex aggregation rule for a mapping to a non-multidimensional source, with a complex aggregation rule for a mapping to a multidimensional source. Instead, you need to create one logical measure for each source, and then create a third logical measure that derives from the first two.
- You can mix noncomplex measures from non-multidimensional data sources with noncomplex measures from multidimensional data sources if they are aggregated through the Oracle BI Server.

Working with Physical Dimensions and Physical Hierarchies

Most dimensions and hierarchies are imported into the Physical layer from multidimensional data sources, rather than created manually. If a particular hierarchy is not imported, any columns associated with that hierarchy are also not imported. If users need access to columns that are not imported, first add these columns to the Physical layer by manually creating them, and then associate them with a level in a hierarchy.

Each level in a hierarchy has a level key. The first cube column associated with (added to) the level of a hierarchy is the level key. This must match with the data source

definition of the cube. The icon for the column that you select first changes to the key icon after it is associated with the level of a hierarchy.

Oracle Business Intelligence supports unbalanced hierarchies for all multidimensional data sources. In general, you can configure unbalanced hierarchies in the Physical layer by changing the hierarchy type.

You can view and edit properties for physical dimensions and hierarchies by double-clicking physical dimension and physical hierarchy objects in the Physical layer of the Administration Tool. You can also view and edit these objects from the Dimensions and Hierarchies tabs of the Cube Table dialog.

This section contains the following topics:

- [Working with Physical Dimension Objects](#)
- [Working with Physical Hierarchy Objects](#)

Working with Physical Dimension Objects

In the Physical Dimension dialog, you can view and edit the name and description of the dimension. You can also add, remove, or edit hierarchies for that dimension, as well as add, remove, or edit columns that represent dimension properties.

Working with Physical Hierarchy Objects

When you select columns to add to a hierarchy, it is recommended that you select them in hierarchical order, starting with the highest level. If you select multiple columns and bring them into the hierarchy at the same time, the order of the selected group of columns remains the same. After adding columns to the hierarchy, you can change the order of the columns in the Browse dialog.

In the Physical Hierarchy dialog, you can view and edit the name and description of the hierarchy, along with the properties described in [Table 7-4](#). For level-based hierarchies, you can add, remove, edit, or reorder levels. For value-based hierarchies, click the Column tab to add, remove, or edit columns. To specify a key column, double-click a column name.

In the Physical Level dialog, you can view and edit the name, external name, and description of the level. You can also add, remove, or edit columns for that level. To designate a column as a level key, double-click a column name.

You should always review the hierarchy type after import to ensure that it is set appropriately. The way this parameter is set upon import depends on the data source. For example, all Essbase hierarchies are initially set to Unbalanced. Review the hierarchy type for each hierarchy and change it as appropriate.

Typically, you always need to manually set the hierarchy type for parent-child (value) hierarchies, except for Hyperion Financial Management hierarchies, which are always set to Value by default upon import. Review the hierarchy type and change the type to Value as appropriate. Parent-child (value) hierarchies are those in which a business transaction, or a cube refresh, can change the number of levels.

For parent-child hierarchies, you must manually set the physical hierarchy type to Value *before* dragging metadata to the Business Model and Mapping layer. The hierarchy type in the Business Model and Mapping layer is set automatically based on the physical hierarchy setting. For all other types, you can determine the hierarchy type later, without needing to rebuild the logical model.

You must also ensure that the corresponding logical dimension properties are correct for queries to work. See [Chapter 9, "Working with Logical Dimensions"](#) for more information.

For SAP/BW data sources, all hierarchies default to fully balanced hierarchies on import. The hierarchy type for two-level hierarchies (which typically correspond to characteristic primary hierarchies) should not be changed. Review all SAP/BW multi-level (external) hierarchies to determine whether any are parent-child hierarchies, and set them to Value as needed.

Table 7–4 Options in the Physical Hierarchy Dialog

Property	Description
External Name	The physical name that is used when referencing the hierarchy in physical MDX queries. This value must reflect the external name defined in the data source.
Dimension Name	(Dimension Unique Name) Dimension to which the hierarchy belongs.
Dimension Type	Identifies whether this hierarchy belongs to a time dimension, measure dimension, or other type of dimension.
Hierarchy Type	<p>Identifies the type of hierarchy, as follows:</p> <ul style="list-style-type: none"> ▪ Fully balanced: A level-based hierarchy with no unbalanced or skip characteristics. Corresponds to a level-based hierarchy in the Business Model and Mapping layer. ▪ Unbalanced: Also called ragged. A hierarchy where the leaves (members with no children) do not necessarily have the same depth. Corresponds to a level-based hierarchy with the Ragged option selected in the Business Model and Mapping layer. ▪ Ragged balanced: Also called skip. A hierarchy where there are members that do not have a value for a particular ancestor level. Corresponds to a level-based hierarchy with the Skipped Levels option selected in the Business Model and Mapping layer. ▪ Network: This hierarchy type is not used. ▪ Value: Also called parent-child. A hierarchy of members that all have the same type. This contrasts with level-based hierarchies, where members of the same type occur only at a single level of the hierarchy. Corresponds to a parent-child hierarchy in the Business Model and Mapping layer. <p>Note: For level-based hierarchies with both unbalanced and skip-level characteristics, choose either Unbalanced or Ragged balanced as the physical hierarchy type. Then, ensure that both Ragged and Skipped Levels are selected for the corresponding logical dimension in the Business Model and Mapping layer.</p>
Default member type ALL	This option is not used.
Use unqualified member name for better performance	Select this option when member names (including aliases) are unique in a given hierarchy so that the Oracle BI Server can take advantage of specific MDX syntax to optimize performance.

Adding or Removing Cube Columns in a Hierarchy After importing a hierarchy, you may need to add or remove a column. If you remove a cube column from a hierarchy, it is deleted from the hierarchy but remains in the cube table and is available for selection to add to other levels.

To add or remove a cube column in an existing hierarchy:

1. In the Physical layer of the Administration Tool, double-click the physical hierarchy that you want to change. The Physical Hierarchy dialog appears.

2. For level-based hierarchies, double-click the level for which you want to add or remove columns. Then, in the Physical Level dialog, you can add, remove, or edit columns. When you are finished, click **OK** in the Physical Level dialog.
3. For value-based hierarchies, click the Columns tab. You can add, remove, or edit columns, as well as designate member key and parent key columns.
4. Click **OK** in the Hierarchy dialog.

Working with Cube Variables for SAP/BW Data Sources

In SAP/BW data sources, cube variables are used as a means of parameterizing queries. Cube variable objects are imported into the Physical layer when metadata is imported from Querycubes/Bex Queries in SAP/BW data sources. Typically, you do not edit these objects directly except to keep them synchronized with the Bex queries in the data source, and except to specify overrides for key characteristics values.

The Cube Variables tab of the Cube Table dialog lists the cube variables for the given cube table, along with the cube variable caption. Double-click a cube variable to see more detailed information, or click the **Add** button to define a new cube variable.

[Table 7-5](#) describes the properties of cube variables for SAP/BW data sources. See your SAP/BW documentation for additional information.

Table 7-5 *Cube Variable Properties*

Property	Description
Name	Name of the cube variable.
Caption	A description (label or caption) associated with the cube variable, mainly used for display purposes.
Variable Type	The type of cube variable. Variable types include: <ul style="list-style-type: none"> ■ SAP_VAR_TYPE_MEMBER: A placeholder for a selection for MEMBER_UNIQUE_NAMES. ■ SAP_VAR_TYPE_HIERARCHY: A placeholder for a HIERARCHY_UNIQUE_NAME. ■ SAP_VAR_TYPE_NUMERIC: A placeholder for a numeric value in formulas.
Selection Type	The selection type of the cube variable, for cube variables of type SAP_VAR_TYPE_MEMBER. Selection types include: <ul style="list-style-type: none"> ■ SAP_VAR_SEL_TYPE_VALUE: The variable is replaced by a single value. Cube variables of type NUMERIC must have this selection type. ■ SAP_VAR_SEL_TYPE_INTERVAL: A placeholder for an interval. ■ SAP_VAR_SEL_TYPE_COMPLEX: A placeholder for a complex selection.

Table 7–5 (Cont.) Cube Variable Properties

Property	Description
Entry Type	Indicates whether replacing variables is optional or mandatory. Entry types include: <ul style="list-style-type: none"> ▪ SAP_VAR_INPUT_TYPE_OPTIONAL: Specifying a value is optional for this variable. ▪ SAP_VAR_INPUT_TYPE_MANDATORY: You must specify a value for this variable. ▪ SAP_VAR_INPUT_TYPE_MANDATORY_NOT_INITIAL: You must specify a value for this variable. An initial field is not a valid entry.
Reference Dimension	This column contains a DIMENSION_UNIQUE_NAME for the parameter type SAP_VAR_TYPE_HIERARCHY.
Reference Hierarchy	This column contains a HIERARCHY_UNIQUE_NAME for the variable type SAP_VAR_TYPE_MEMBER.
Default Low	This property contains a default value for the variable or is zero.
Default High	This property contains a default value for the variable or is zero. This property is only important for variables with the selection type SAP_VAR_SEL_TYPE_INTERVAL and SAP_VAR_SEL_TYPE_SELECTION.
Override Default Low	Provide a default value for the cube variable in this field if the Default Low is zero. You must specify a value for this property for mandatory variables that do not specify a default value.
Override Default High	Provide a default value for the cube variable in this field if the Default High is zero. You must specify a value for this property for mandatory variables that do not specify a default value.

Viewing Members in Physical Cube Tables

You can view members of hierarchies or levels in the Physical layer of repositories. To view members, the repository must be open in online mode. The list of members by level in the hierarchy can help you determine if the connection pool is set up properly. You might want to reduce the time it takes to return data or the size of the returned data by specifying a starting point (Starting from option) and the number of rows you want returned (Show option).

To view members:

1. Open the Administration Tool in online mode.
2. In the Physical layer, right-click a hierarchy or level.
3. Select **View Members**.

A window opens showing the number of members in the hierarchy and a list of the levels. You might need to enlarge the window and the columns to view all the returned data.

4. Click **Query** to display results.
5. When finished, click **Close**.

Working with Essbase Data Sources

This section describes how Essbase data is modeled by default in the Physical layer of the Oracle BI repository, and describes the tasks you can perform to model the data in different ways.

This section contains the following topics:

- [About Using Essbase Data Sources with Oracle Business Intelligence](#)
- [Working with Essbase Alias Tables](#)
- [Associating Member Attributes to Dimensions and Levels](#)
- [Modeling Alternate Hierarchies](#)
- [Modeling Measure Hierarchies](#)
- [Improving Performance by Using Unqualified Member Names](#)

About Using Essbase Data Sources with Oracle Business Intelligence

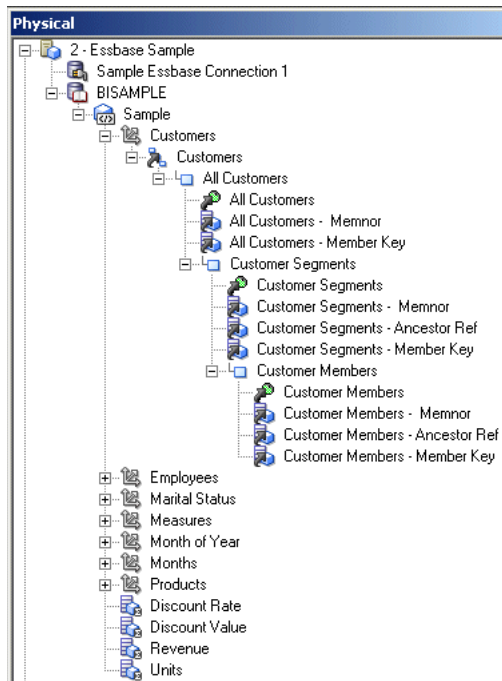
When you import metadata from Essbase data sources, the cube metadata is mapped to the Physical layer in a way that supports the Oracle Business Intelligence logical model. Metadata that applies to all members of the dimension, such as aliases, are modeled as dimension properties by default. Level-based properties, such as outline sort/memnor information, are mapped as separate physical cube columns in the dimension.

The following physical column types are used for Essbase metadata:

- **Member Alias:** Indicates an Alias column.
- **UDA:** Indicates the column is a User Defined Attribute (UDA).
- **Outline Sort:** Indicates the column is of memnor type, used for outline sorts in the logical layer. Imported at the lowest level of each dimension.
- **Attribute:** Indicates the column is of attribute type, for attribute dimensions.
- **Other:** The type is different than those listed, or unknown.
- **Ancestor Reference:** References the ancestor of a dimension.
- **Member Key:** Indicates the column is a member key.
- **Leaf:** Indicates that the column is the lowest member of the hierarchy.
- **Root:** Indicates that the column is the root member of the hierarchy.
- **Parent Reference:** References the parent of a dimension.

The column types Outline Sort, Ancestor Reference, Member Key, Leaf, Root, and Parent Reference are used internally by the system and should not be changed.

[Figure 7–2](#) shows Essbase data that has been imported into the Physical layer.

Figure 7-2 Essbase Data Modeled in the Physical Layer

There are different options in the Physical layer that let you control how you want to model certain types of metadata. Choose the option that best meets the needs of your user base. For example, many types of Essbase metadata are modeled as dimension properties by default in the Physical layer. This multidimensional structure works best with the new hierarchical reporting style introduced in the current release.

Alternatively, you can choose to flatten the Essbase metadata in the Physical layer for ease of use with the attribute-style reporting supported in previous releases of Oracle Business Intelligence.

The following list summarizes some of these modeling options:

- **Aliases.** Aliases are modeled as dimension properties by default, but you can also choose to flatten them using the **Create Columns for Alias Table** feature. See ["Working with Essbase Alias Tables"](#) for more information.
- **UDAs.** UDAs are modeled as dimension properties by default, but you can also choose to flatten them using the **Create Columns for UDA** feature. See ["Modeling User-Defined Attributes"](#) for more information.
- **Alternate Hierarchies.** Alternate hierarchies are modeled as separate hierarchies by default, but you can choose to view them in as a single hierarchy using the **Convert to single hierarchy view** feature. See ["Modeling Alternate Hierarchies"](#) for more information.
- **Measure Hierarchies.** By default, measures are imported as a single measure column that represents all the measures, but you can also choose to view each measure as an individual column using the **Convert measure dimension to flat measures** feature. See ["Modeling Measure Hierarchies"](#) for more information.

Note the following additional information about using Essbase data sources with Oracle Business Intelligence:

- **Substitution variables.** Essbase substitution variables are automatically retrieved and populated into corresponding Oracle BI Server session variables. Depending

on the scope of the Essbase variable, the naming convention for the Oracle BI Server variable is as follows:

Server instance scope: *server_name:var_name*

Application scope: *server_name:app_name:var_name*

Cube scope: *server_name:app_name:cube_name:var_name*

A single initialization block is also created in the repository for the Essbase variables. Set the appropriate refresh interval in the initialization block to reflect anticipated update cycles for Essbase variables.

- **Essbase Generations.** Essbase Generations are mapped to physical level objects.
- **Time series functions.** The Oracle BI Server time series functions AGO, TODATE, and PERIODROLLING are pushed down to Essbase to take advantage of the native capabilities of the Essbase server.
- **Database functions.** You can use the database SQL functions EVALUATE and EVALUATE_AGGREGATE to leverage functions specific to Essbase data sources. See ["Examples Using EVALUATE_AGGREGATE and EVALUATE to Leverage Unique Essbase Functions"](#) for more information.

Note that EVALUATE_PREDICATE is not supported for use with Essbase data sources.

- **Gen 1 levels.** By default, Gen 1 levels are included when you drag and drop an Essbase cube or dimension from the Physical layer to the Business Model and Mapping layer. However, because Gen 1 levels are not usually needed for analysis, you can choose to exclude Gen 1 levels when you drag and drop Essbase objects to your business model. To do this, select **Skip Gen 1 levels in Essbase drag and drop actions** in the General tab of the Options dialog. See ["Setting Preferences"](#) for more information.
- **Hierarchy types.** For Essbase data sources, all hierarchies are imported as Unbalanced by default. Review the **Hierarchy Type** property for each physical hierarchy and change the value if necessary. Supported hierarchy types for Essbase are **Unbalanced**, **Fully balanced**, and **Value**.

About Incremental Import

You can choose to incrementally import Essbase metadata. In other words, you can perform an initial import, and then import again. You might want to import incrementally when information in the data source has changed, or when your first import only included a subset of the metadata. Note the following about incremental import:

- When you re-import metadata that already exists in the Physical layer, a message appears, warning you that your Physical objects will be overwritten.
- If you delete data in the source, re-importing the metadata does not automatically perform the deletion in the Physical layer. Instead, you must manually delete the corresponding Physical objects.
- If you rename an object in the source, the renamed object is imported as a new object. In this case, both the old object and the new (renamed) object are displayed in the Physical layer.
- In general, customizations that you have performed on the Physical layer data, such as determining the alias column to use for display, are retained after an incremental import. If you want to revert to the default imported view, you must delete the existing Physical layer objects and then re-import the metadata.

Working with Essbase Alias Tables

Essbase cubes support the concept of aliases, which are alternate names for members or shared members. For example, the member name might be a product code (100), with a default alias for the product name (Cola) and an additional alias for the long name (Cherry Cola). Often, members have separate aliases for each user language to enable users to view member names in their own language.

In the Essbase cube, aliases are stored in alias tables that map a specific set of alias names to member names. Typically, a Default alias table exists for each cube.

This section contains the following topics:

- [Determining the Value to Use for Display](#)
- [Explicitly Defining Columns for Each Alias](#)

Determining the Value to Use for Display

When you import metadata from Essbase into the Oracle BI repository, the Essbase cube table object in the Physical layer has a property that determines which value to display for members: the member name, the default alias name, or some other alias name. By default, the columns display the default alias name.

To change the value to display for members:

1. In the Physical layer of the Administration Tool, double-click an Essbase cube table.
2. In the General tab of the Cube Table dialog, choose the appropriate value for **Display Column**. You can select **Member Name**, or you can select **Alias** and then choose an alias table name from the list.
3. Click **OK**.

Explicitly Defining Columns for Each Alias

Aliases are modeled as dimension properties in the Physical layer after import. If you want to work with more than one alias, such as when you want to flatten attributes for reporting purposes or externalize strings for translation, you can explicitly define columns for each alias. You can define alias columns at the cube, dimension, or hierarchy level.

To explicitly define columns for each alias:

1. In the Administration Tool, in the Physical layer, right-click the cube table, physical dimension, or physical hierarchy for which you want to define alias columns.
2. Select **Create Columns for Alias Table**. Then, from the sub-list, select the alias table for which you want to create columns.

Note that the **Fetch** button is not used.

3. Click **Create**.
4. Drag the new alias columns to the appropriate location in the Business Model and Mapping layer.

If you want to externalize strings for translation based on the alias columns, see "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Modeling User-Defined Attributes

Essbase supports the concept of user-defined attributes (UDAs). A UDA is essentially any arbitrary textual string that can be associated with any member from a dimension. A member can have multiple strings associated to it.

You can choose whether to import UDAs in the Import Metadata Wizard. If you choose to import UDAs, then by default, each UDA is modeled as a dimension property in the Physical layer of the repository.

You can also choose to model each UDA as a separate physical column. To do this, perform one of the following tasks:

- To model all UDAs in a cube as separate physical columns, right-click the cube table and select **Create columns for UDA**. All UDAs in the cube are modeled as separate physical columns.
- To model all UDAs in a dimension as separate physical columns, right-click the dimension object and select **Create columns for UDA**, then select **All UDAs**. All UDAs in the dimension are modeled as separate physical columns.
- To model a particular UDA in a dimension as a separate physical column in each level, right-click the dimension object and select **Create columns for UDA**, then select the specific UDA you want to model. The selected UDA is modeled as a separate physical column for each level.

Associating Member Attributes to Dimensions and Levels

Member attributes are not automatically associated to corresponding dimensions and levels during the import process. To manually create the association, map the member attribute to the appropriate logical table. In other words, drag and drop the columns from the attribute dimension in the Physical layer to the appropriate logical tables in the Business Model and Mapping layer.

Modeling Alternate Hierarchies

By default, alternate hierarchies are modeled as separate hierarchies in the Physical layer. You can choose to view them as separate hierarchies (called the multi-hierarchy view), or as a single hierarchy.

To view alternate hierarchies as a single hierarchy, right-click the dimension object containing the alternate hierarchies and select **Convert to single hierarchy view**. To return to the multi-hierarchy view, right-click the dimension object again and select **Convert to multi-hierarchy view**.

For example, [Figure 7-3](#) shows the multi-hierarchy view for an alternate hierarchy.

Figure 7-3 Essbase Alternate Hierarchy Displayed in Multi-Hierarchy View

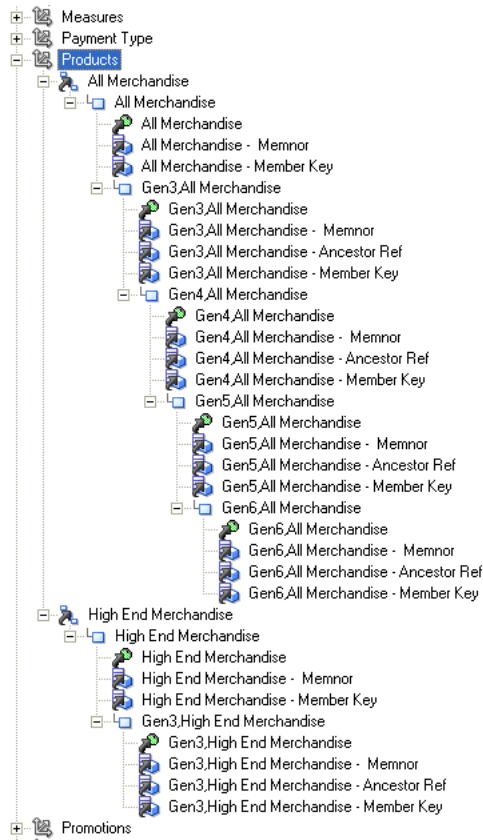



Figure 7-4 shows the single-hierarchy view for the same alternate hierarchy.

Figure 7–4 Essbase Alternate Hierarchy Displayed in Single-Hierarchy View

Modeling Measure Hierarchies

By default, measures are imported as measure hierarchies. In other words, the cube contains a single measure column that represents all the measures.

Alternatively, you can choose to flatten the measure hierarchy to view each measure as an individual column. To do this, right-click the cube object and select **Convert measure dimension to flat measures**.

Improving Performance by Using Unqualified Member Names

When member names (including aliases) are unique in a given hierarchy, the Oracle BI Server can take advantage of specific MDX syntax to optimize performance. To enable this capability, select **Use unqualified member name for better performance** in the Hierarchy dialog.

The import process cannot identify that member names are unique for a given hierarchy, so it is the responsibility of the administrator to confirm uniqueness. Note that query errors may result if a hierarchy is specified as having unique members when it does not.

Note: If you find that the Oracle BI Server is generating incorrect queries for Essbase, check to see whether there are duplicate member names in a given hierarchy. If there are, ensure that the option **Use unqualified member name for better performance** is not selected for that hierarchy, or perform the following steps:

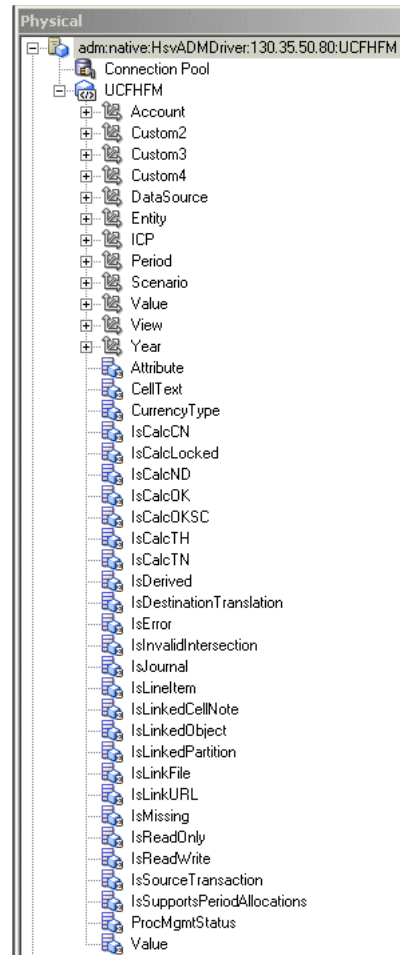
1. From the Essbase outline, update each offending member variable by adding a prefix or suffix to make the member name unique.
 2. Update SQL queries as necessary, if references are made to data within SQL.
 3. Reload the data and members in the Essbase outline.
-
-

Working with Hyperion Financial Management Data Sources

When you import data from Hyperion Financial Management data sources, both measures and dimensions are imported into the Physical layer. The Hyperion Financial Management hypercube model is exposed in the Physical layer in the following ways:

- There is only one measure in Hyperion Financial Management, called Value. This measure is modeled as a single fact column in the Physical layer, also called Value.
- The Value measure has three base properties: CellText, CurrencyType, and Attribute. These properties are all represented as additional fact columns.
- The Attribute property has additional sub-properties, such as IsReadOnly. These properties are also exposed as additional columns.

[Figure 7–5](#) shows how Hyperion Financial Management data is modeled in the Physical layer.

Figure 7–5 Hyperion Financial Management Metadata in the Physical Layer

All Hyperion Financial Management dimensions are modeled as parent-child hierarchies in the Physical layer. Shared members, alternate hierarchies, and unbalanced hierarchies are supported.

Dimension member properties are exposed as columns (such as Name, Description, ShortName, and so on). An additional column called Sort Order is also displayed for each dimension. This column contains custom sort information retrieved from the Hyperion Financial Management data source.

Each Hyperion Financial Management dimension has a corresponding Point of View (POV) value that provides customized information for different users. This POV value is exposed as the **Default Member** in the Hierarchies tab of the Dimension dialog. Although the Default Member field is populated upon import, note that the default values may need to be updated according to the needs of the user base.

Note: Do not select the **Default member type ALL** option for Hyperion Financial Management hierarchies.

About Query Support for Hyperion Financial Management Data Sources

Both member queries (dimensional browsing) and data queries (measure analysis) are supported for Hyperion Financial Management data sources.

Most Logical SQL functions are performed in the Oracle BI Server. However, you should use `EVALUATE_PREDICATE` to access the following functions specific to Hyperion Financial Management:

- `PeriodOffset` (used to access prior or future periods through an offset)
- NA Suppression functions specific to Hyperion Financial Management (`SuppressDerived`, `SuppressInvalidIntersection`, `SuppressMissing`, `SuppressNoAccess`, `SuppressZero`, `SuppressError`)
- Base function (returns the leaf members below a given ancestor member)
- `CommonChildren`
- User-defined functions

See "[EVALUATE_PREDICATE](#)" for detailed information about syntax and usage.

Note that there is no native support for time series functions. Time series functions are only supported through data modeling.

Working with Oracle OLAP Data Sources

Oracle Database has an OLAP Option that provides an embedded, full-featured online analytical processing server. The OLAP Option is used in the following roles:

- A summary management solution to SQL-based business intelligence tools and applications.
- A calculation engine that provides SQL-based business intelligence tools with rich analytic content.
- A full-featured multidimensional server, servicing dimensionally oriented business intelligence tools and applications.

Oracle Business Intelligence supports Oracle OLAP as a data source. When you import metadata from an Oracle OLAP source, the Oracle OLAP objects appear in the Physical layer of the Administration Tool. This section provides information about the Physical layer Oracle OLAP objects.

See *Oracle OLAP User's Guide* for more information about Oracle OLAP.

This section contains the following topics:

- [About Importing Metadata from Oracle OLAP Data Sources](#)
- [Working with Oracle OLAP Analytic Workspace \(AW\) Objects](#)
- [Working with Oracle OLAP Dimensions, Hierarchies, and Levels](#)
- [Working with Oracle OLAP Cubes and Columns](#)

About Importing Metadata from Oracle OLAP Data Sources

This section provides important information about using the Administration Tool to import metadata from Oracle OLAP, as follows:

- For Oracle OLAP cubes with multi-language metadata, only the default language is imported.
- Only dimensions that contain at least one hierarchy are imported.
- Multiple hierarchies in a single query are not supported. If a query includes columns from multiple hierarchies in a given dimension, the Oracle BI Server returns an error.

- The default aggregation rule in the Business Model and Mapping layer for Oracle OLAP measures is External Aggregation. The External Aggregation rule means that the Oracle BI Server is not aware of the underlying aggregation rule for the specific measure and will not compute it internally. Instead, the Oracle BI Server will always ship the query to the underlying multidimensional data source for aggregation.

In some cases, you may want to set the aggregation rule for a measure to something other than External Aggregation. For example, you may have federated multiple data sources, or you may want to perform higher-level aggregation along dimension attributes that are not represented by a level in Oracle OLAP. In both of these cases, you can change the default aggregation rule to match the rule in the underlying data source or sources. Note that the aggregation is still performed in the Oracle OLAP data source where possible.

See "[System Requirements and Certification](#)" for the latest information about the versions of Oracle OLAP supported by Oracle Business Intelligence.

Working with Oracle OLAP Analytic Workspace (AW) Objects

You can view Oracle OLAP Analytic Workspace (AW) objects in the Physical layer of the Administration Tool. These objects correspond to the analytic workspace object in the Oracle OLAP metadata, and are similar to physical catalog and physical schema objects. Analytic workspaces are containers for storing related cubes. You create dimensions, cubes, and other dimensional objects within the context of an analytic workspace.

Oracle OLAP Analytic Workspace objects have properties for **Name**, **Description**, and **Dynamic Name**. You can use the Dynamic Name tab to provide a variable that specifies the name of the Analytic Workspace object. Note that the Dynamic Name tab is not active unless at least one session variable is defined. See "[Using a Variable to Specify the Name of a Catalog or Schema](#)" for more information.

Working with Oracle OLAP Dimensions, Hierarchies, and Levels

Oracle OLAP dimensions are lists of unique values that identify and categorize data. They form the edges of a cube, and thus of the measures within the cube. In a report, the dimension values (or their descriptive attributes) provide labels for the rows and columns. There are three types of Oracle OLAP dimensions:

- **Level-based dimensions.** Members of level-based dimensions naturally group into levels based on their type, such as 'month' and 'year.' Most dimensions are level-based.
- **Value-based dimensions.** These dimensions have parent-child relationships among their members, but the members are all the same type (like 'Employee' or 'Account'), so these relationships do not form meaningful levels.
- **List or flat dimensions.** These dimensions have no levels or hierarchies.

Note: Oracle Business Intelligence does not support dimensions that have no hierarchies (flat dimensions). Importing flat dimensions from an Oracle OLAP data source will result in an error. If you have flat dimensions, replace them with single-level hierarchies in the data source before importing them into Oracle Business Intelligence.

On the General tab of the Oracle OLAP Dimension dialog, you can view and edit the name and description of the dimension, along with the following dimension properties:

- **Time.** Indicates that this dimension is a time dimension.
- **Ragged.** Indicates that this dimension contains a hierarchy that has at least one member with a different base, creating a "ragged" base level for the hierarchy.
- **Skipped levels.** Indicates that this dimension contains a hierarchy that has at least one member whose parents are more than one level above it, creating a hole in the hierarchy. An example of a skip-level hierarchy is City-State-Country, where at least one city has a country as its parent (for example, Washington D.C. in the United States).
- **External Name.** The physical name that is used when referencing the dimension in physical SQL queries. This value must reflect the external name defined in the data source.
- **Cache properties.** Select **Cacheable** to include this dimension in the Oracle BI Server query cache. To specify that cache entries do not expire, select **Cache never expires**. Alternatively, you can select **Cache persistence time** and enter a value to specify how long entries should persist in the query cache. Note that if a query references multiple physical objects with different persistence times, the cache entry for the query exists for the shortest persistence time set for any of the tables referenced in the query. This makes sure that no subsequent query gets a cache hit from an expired cache entry.

The Columns and Hierarchies tabs of the Oracle OLAP Dimension dialog list the dimension members and hierarchies that belong to the dimension. In the Columns tab, you can add or remove columns, as well as edit particular columns. In the Hierarchies tab, you can add, remove, or edit hierarchies. You can also use the type (key) button to select the default hierarchy for the dimension.

Dimensions can contain one or more hierarchies. Most hierarchies are level-based and consist of one or more levels of aggregation. Members roll up into the next higher level in a many-to-one relationship, and these members roll up into the next higher level, and so forth to the top level. Ragged and skip-level hierarchies are also supported.

Dimensions can also contain value-based hierarchies, which are parent-child hierarchies that do not support levels. For example, an employee dimension might have a parent-child relationship that identifies each employee's supervisor. However, levels that group together first-, second-, and third-level supervisors and so forth may not be meaningful for analysis.

Multiple hierarchies for a dimension typically share the base-level dimension members and then branch into separate hierarchies. They can share the top level if they use all the same base members and use the same aggregation operators. Otherwise, they need different top levels to store different aggregate values.

In the Oracle OLAP Hierarchy dialog, you can view and edit the name, external name, and description of the hierarchy. For level-based hierarchies, you can add, remove, edit, or reorder levels. For value-based hierarchies, you can add, remove, or edit columns. To specify a key column, double-click a column name.

In the Oracle OLAP Level dialog, you can view and edit the name, external name, and description of the level. You can also add, remove, or edit columns for that level. To designate a column as a level key, double-click a column name.

Working with Oracle OLAP Cubes and Columns

Oracle OLAP cubes are informational objects that identify measures with the exact same dimensions and thus are candidates for being processed together at all stages: data loading, aggregation, storage, and querying. Cubes define the shape of your business measures. They are defined by a set of ordered dimensions. The dimensions form the edges of a cube, and the measures are the cells in the body of the cube.

Oracle OLAP cubes have properties similar to other cubes. On the General tab of the Oracle OLAP Cube dialog, you can view and edit the name and description of the cube, along with the following cube properties:

- **External Name.** The physical name that is used when referencing the cube in physical SQL queries. This value must reflect the external name defined in the data source.
- **Density and Materialization.** For Oracle OLAP 10g cubes that are sparse and fully materialized, you should specify values for these properties to optimize queries. If you set the **Density** option to **Sparse** and the **Materialization** option to **Fully Materialized**, the Oracle BI Server generates a loop clause to skip empty cells. Note that if you leave the **Density** option blank, the Oracle BI Server assumes the data is sparse.

If you set these options, make sure that you set them to reflect the actual properties of the data source. Do not specify that the data is sparse and fully materialized unless this is true for your data source.

You do not need to set these values for Oracle OLAP 11g cubes. For these objects, optimization happens automatically.

- **Cache properties.** Select **Cacheable** to include this cube in the Oracle BI Server query cache. To specify that cache entries do not expire, select **Cache never expires**. Alternatively, you can select **Cache persistence time** and enter a value to specify how long entries should persist in the query cache. Note that if a query references multiple physical objects with different persistence times, the cache entry for the query exists for the shortest persistence time set for any of the tables referenced in the query. This makes sure that no subsequent query gets a cache hit from an expired cache entry.

The Columns tab of the Oracle OLAP Cube dialog lists the columns that belong to the cube. You can add or remove columns, as well as edit particular columns.

Oracle OLAP columns can be measures, calculated measures, attributes, or level keys. Oracle OLAP columns have the same properties as other physical columns. See ["Creating and Editing a Column in a Physical Table"](#) for more information about physical column properties like **Type**, **Length**, and **Nullable**.

Working with Physical Foreign Keys and Joins

You can create physical foreign keys and complex joins using either the Physical Diagram, or the Joins Manager. Note that you do not create joins for multidimensional data sources.

This section contains the following topics:

- [About Physical Joins](#)
- [Defining Physical Joins with the Physical Diagram](#)
- [Defining Physical Joins with the Joins Manager](#)

About Physical Joins

All valid physical joins need to be configured in the Physical layer of the Administration Tool.

When you import keys in a physical schema, the primary key-foreign key joins are automatically defined. Any other joins within each data source or between data sources have to be explicitly defined to express relationships between tables in the Physical layer.

Imported key and foreign key joins do not have to be used in metadata. Joins that are defined to enforce referential integrity constraints can result in incorrect joins being specified in queries. For example, joins between a multipurpose lookup table and several other tables can result in unnecessary or invalid circular joins in the SQL queries issued by the Oracle BI Server.

This section contains the following topics:

- [About Primary Key and Foreign Key Relationships](#)
- [About Complex Joins](#)
- [About Multi-Database Joins](#)
- [About Fragmented Data](#)

About Primary Key and Foreign Key Relationships

A primary key and foreign key relationship defines a one-to-many relationship between two tables. A foreign key is a column or a set of columns in one table that references the primary key columns in another table. The primary key is defined as a column or set of columns where each value is unique and identifies a single row of the table.

Note that there are two cases where multiple foreign key columns in a table point to the same table:

- When the primary key of the foreign table is "concatenated," meaning that it consists of a set of columns. This is a single join between two tables that happens to use multiple columns.
- When you have created an alias to the foreign table, because the foreign table needs to serve in different roles. In this case, each foreign key joins to a primary key in one role-playing alias or the other. See "[About Physical Alias Tables](#)" for more information.

You can specify primary key and foreign keys in the Physical Diagram, or by using the Keys and Foreign Keys tabs of the Physical Table dialog. Also refer to "[Defining Physical Joins with the Physical Diagram](#)" and "[Creating and Managing Columns and Keys for Relational and Cube Tables](#)" for more information.

About Complex Joins

In the Physical layer of the repository, complex joins are joins over nonforeign key and primary key columns. In other words, physical complex joins are joins that use an expression rather than key column relationships. When you create a complex join in the Physical layer, you specify the expression for the join.

For most data sources, foreign key joins are preferred for performance reasons. Complex joins are usually not as performant because they do not use key column relationships to form the join. The exception is ADF business component data sources, which use physical complex joins exclusively to denote ViewLink instances that connect pairs of View Objects in the ADF model.

About Multi-Database Joins

A multi-database join is defined as a table under one metadata database object that joins to a table under a different metadata database object. You need to specify multi-database joins to combine the data from different databases. Use the Physical Diagram to specify multi-database joins. See ["Defining Physical Joins with the Physical Diagram"](#) for more information.

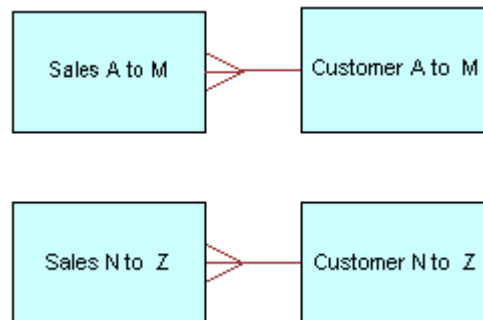
Multi-database joins can be created between tables in most types of databases and are performed within the Oracle BI Server. Note that you cannot create multi-database joins to tables in Oracle OLAP data sources.

While the Oracle BI Server has several strategies for optimizing the performance of multi-database joins, these joins are significantly slower than joins between tables within the same database. For this reason, you should avoid them whenever possible.


About Fragmented Data

Fragmented data is data from a single domain that is split between multiple tables. For example, a data source might store sales data for customers with last names beginning with the letter A through M in one table and last names from N through Z in another table. With fragmented tables, you need to define all of the join conditions between each fragment and all the tables to which it relates. [Figure 7-6](#) shows the physical joins with a fragmented sales table and a fragmented customer table where they are fragmented the same way (A through M and N through Z).

Figure 7-6 *Fragmented Tables Example*



In some cases, you might have a fragmented fact table and a fragmented dimension table, but the fragments might be across different values. In this case, in addition to the joins created in [Figure 7-6](#), you need to define a one-to-many join from Customer A to F and from Customer G to Z to Sales A to M, as shown in [Figure 7-7](#).

Figure 7-7 Joins for Fragmented Tables Example

Note: Avoid adding join conditions where they are not necessary (for example, between Sales A to M and Customer N to Z in [Figure 7-6](#)). Extra join conditions can cause performance degradations.

Defining Physical Joins with the Physical Diagram

You can define foreign keys and complex joins between tables, whether or not the tables are in the same data source. When you use the Physical Diagram to create joins, the Administration Tool determines what type of join to create based on the selected object types and the join expression.

If you do not want the Administration Tool to automatically determine what type of join to create, use the Joins manager to explicitly create the join. See "[Defining Physical Joins with the Joins Manager](#)" for more information.

To define a physical foreign key join or a complex join with the Physical Diagram:

1. In the Physical layer of the Administration Tool, select one or more tables and choose one of the Physical Diagram commands from the right-click menu.
2. Click the **New Join** button on the Administration Tool toolbar:



3. In the Physical Diagram, left-click the first table in the join (the table representing many in the one-to-many join) to select it.
4. Move the cursor to the table to which you want to join (the table representing one in the one-to-many join), and then left-click the second table to select it.

The Physical Foreign Key dialog appears. Although physical foreign key joins are the default join type, the object type might change to a complex join after you define the join and click **OK**, depending on the join information.

5. Select the joining columns from the left and the right tables.

The SQL join conditions appear in the expression pane.

The driving table option is shown in this dialog, but it is not available for selection because the Oracle BI Server implements driving tables only in the Business Model and Mapping layer. See "[Specifying a Driving Table](#)" for more information about driving tables.

6. For complex joins, you can optionally set the cardinality for each side of the join (for example, **N**, **0,1**, **1**, or **Unknown**).

To set the cardinality to unknown, you only need to select **Unknown** for one side of the join. For example, choosing unknown-to-1 is equivalent to unknown-to-unknown and appears as such the next time you open the dialog for this join.


7. If appropriate, specify a database hint. See ["Using Hints"](#) for more information.
8. If you are creating a complex join for ADF Business Component ViewObject or ViewLink instances, specify the ViewLink instance name or the ViewLink definition name in the **ViewLink Name** field.
9. To open Expression Builder, click the button to the right of the Expression pane. The expression displays in the Expression pane.

The default join expression for ViewObject or ViewLink instances is arbitrary and has no meaning.

10. Click **OK** to apply the selections.

In the Physical Diagram, the join is represented by a line between the two selected tables, with an arrow at the "one" end of the join. [Figure 7-8](#) shows a join in the Physical Diagram.

Figure 7-8 Join in the Physical Diagram



Defining Physical Joins with the Joins Manager

You can use the Joins Manager to view join relationships and to create physical foreign key joins and complex joins.

To define a physical foreign key join or complex join with the Joins Manager:

1. In the Administration Tool toolbar, select **Manage**, then select **Joins**.
2. In the Joins Manager dialog, perform one of the following tasks:
 - Select **Action > New > Complex Join**.
The Complex Join dialog appears.
 - Select **Action > New > Physical Foreign Key**. Then, in the Browse dialog, double-click a table.
3. In the Complex Join or Physical Foreign Key dialog, type a name for the join.
4. Click the **Browse** button for the **Table** field on the left side of the dialog, and then locate the table that the foreign key references.
5. Select the columns in the left table that the key references.
6. Select the columns in the right table that make up the foreign key columns.
7. For complex joins, you can optionally set the cardinality for each side of the join (for example, **N**, **0,1**, **1**, or **Unknown**).

To set the cardinality to unknown, you only need to select **Unknown** for one side of the join. For example, choosing unknown-to-1 is equivalent to unknown-to-unknown and appears as such the next time you open the dialog for this join.

8. If appropriate, specify a database hint. See ["Using Hints"](#) for more information.

9. If you are creating a complex join for ADF Business Component ViewObject or ViewLink instances, specify the ViewLink instance name or the ViewLink definition name in the **ViewLink Name** field.
10. To open Expression Builder, click the button to the right of the Expression pane. The expression displays in the Expression pane.

The default join expression for ViewObject or ViewLink instances is arbitrary and has no meaning.
11. Click **OK** to save your work.

Deploying Opaque Views

An opaque view is a Physical layer table that consists of a `SELECT` statement. When you need a new table, you should create a physical table or a materialized view. An opaque view should be used only if there is no other solution. See [Appendix F, "Exchanging Metadata with Databases to Enhance Query Performance"](#) for more information about materialized views.

This section contains the following topics:

- [About Deploying Opaque Views](#)
- [Deploying Opaque View Objects](#)
- [Undeploying a Deployed View](#)
- [When to Delete Opaque Views or Deployed Views](#)
- [When to Redeploy Opaque Views](#)

About Deploying Opaque Views

In the repository, opaque views appear as view tables in the data source, but the view does not actually exist until you deploy it. You deploy an opaque view in the data source using the Deploy Views utility. After deploying an opaque view, it is called a deployed view. Opaque views can be used without deploying them, but the Oracle BI Server has to generate a more complex query when an opaque view is encountered.

Note: Data sources such as XLS and nonrelational data sources do not support opaque views and cannot run the view deployment utility.

To verify that opaque views are supported by a data source, check whether the `CREATE_VIEW_SUPPORTED` SQL feature is selected in the Database dialog, in the Features tab. See ["Specifying SQL Features Supported by a Data Source"](#) for instructions.

Deploying Opaque View Objects

In offline mode, the Deploy Views utility is available when importing from data sources with ODBC and DB2 CLI data sources. Oracle Native (client) drivers are also supported in the offline mode for deploying views. In online mode, view deployment is available for supported data sources using Import through server (the settings on the client are ignored).

Using the Create View SELECT Statement

The SQL statement for deploying opaque views in the Physical layer of the repository is available for supported data sources. To determine which of your data sources support opaque views, contact your system administrator or consult your data source documentation.

Only repository variables can be used in the definition. An error is generated if a session variable is used in the view definition.

Syntax

```
CREATE VIEW view_name AS select_statement,
```

Where:

- *select_statement* is the user-entered SQL in the opaque view object. If SQL is invalid, the create view statement fails during view deployment.
- *view_name* is one of the two following formats: *schema.viewname*, or *viewname*. The connection pool settings determine if the schema name is added.

For opaque view objects, the right-click menu contains the **Deploy View(s)** option. When you select **Deploy View(s)**, the Create View SQL statement executes and attempts to create the deployed view objects. The following list describes the ways you can initiate view deployment and the results of each method:

- Right-click a single opaque view object. When you select **Deploy View(s)**, the Create View SQL statement executes and attempts to create a deployed view for the object.
- Right-click several objects. If at least one of the selected objects is an opaque view object, the right-click menu contains the **Deploy View(s)** option. When you select **Deploy View(s)**, the Create View SQL statement executes and attempts to create the deployed views for any qualifying objects.
- Right-click a physical schema or physical catalog. If any opaque view object exists in the schema or catalog, the right-click menu contains the **Deploy View(s)** option. When you select **Deploy View(s)**, the Create View SQL statements for all qualifying objects execute and attempt to create deployed views for the qualifying objects contained in the selected schema or catalog.

During deployment, names are assigned to the views. If you change the preassigned name, the new name must be alphanumeric and no more than 18 characters. If these guidelines are not followed, the object name is automatically transformed to a valid name using the following Name Transform algorithm:

1. All non-alphanumeric characters are removed.
2. If there are 16 or more characters after Step 1, the first 16 characters are kept.
3. Two digits starting from 00 to 99 are appended to the name to make the name unique in the corresponding context.

After the deployment process completes, the following occurs:

- Views that have been successfully and unsuccessfully deployed appear in a list.
- For unsuccessful deployments, a brief reason appears in the list.
- If deployment is successful, the object type of the opaque view changes from Select to None and the deployed view is treated as a regular table.

If you change the type back to Select, the associated opaque views are dropped from the data source, or an error message appears. See ["When to Delete Opaque Views or Deployed Views"](#) for information about deleting deployed views.

- In the Administration Tool, the view icon changes to the deployed view icon for successfully deployed views.

To deploy an opaque view:

1. In the Physical layer of the Administration Tool, right-click the opaque view that you want to deploy.
2. In the right-click menu, select **Deploy View(s)**.
3. In the View Deployment - Deploy View(s) dialog, perform the following steps:
 - a. In the **New Table Name** column, you can optionally change the new deployed view names.

If the change does not conform to the naming rules, a new name is assigned and the dialog appears again so that you can accept or change it. This action repeats until all names pass validation.
 - b. If you do not want to deploy one or more of the views, clear the appropriate rows.
4. If there are multiple connection pools defined for the physical database, in the Select Connection Pool dialog, choose a connection pool and click **Select**.

The SQL statement (`CREATE VIEW`) executes, and the View Deployment Messages dialog appears.
5. In the View Deployment Messages dialog, you can search for views using **Find** and **Find Again**, or copy the contents.
6. When you are finished, click **OK**.

Undeploying a Deployed View

Running the Undeploy Views utility against a deployed view deletes the views and converts the table back to an opaque view with its original `SELECT` statement.

To undeploy a deployed view:

1. In the Physical layer of the Administration Tool, right-click a physical database, catalog, schema, or table.

If a deployed view exists that is related to the selected object, the right-click menu contains the **Undeploy View(s)** option.
2. Select **Undeploy View(s)**.

A list of views to be undeployed appears.
3. If you do not want to undeploy one or more of the views, clear the appropriate rows.
4. In the View Deployment - Undeploy View(s) dialog, click **OK** to remove the views.

A message appears if the undeployment was successful.
5. In the View Deployment Messages dialog, you can search for undeployed views using **Find** and **Find Again**, or you can copy the contents.
6. When you are finished, click **OK**.

When to Delete Opaque Views or Deployed Views

Use the following guidelines to remove opaque or deployed view objects in the repository:

- **Removing an undeployed opaque view in the repository.** If the opaque view has not been deployed, you can delete it from the repository.
- **Removing a deployed view.** When you deploy an opaque view, a view table is created physically in both the data source and the repository. Therefore, you must undeploy the view before deleting it. You use the Undeploy Views utility in the Administration Tool. This removes the opaque view from the back-end data source, changes the Table Type from None to Select, and restores the `SELECT` statement of the object in the Physical layer of repository.

Caution: Do not manually delete the view table in the data source. If deleted, the Oracle BI Server cannot query the view object. When you undeploy the view, it is removed automatically from the data source.

When to Redeploy Opaque Views

After removing an opaque view, you can choose to redeploy it. The Administration Tool does not distinguish between a first-time deployment and a redeployment. Make sure that you remove a deployed view before deploying the opaque view again. Failure to do this causes the deploy operation to fail, and an error message is returned from the data source.

Using Hints

Hints are instructions placed within a SQL statement that tell the data source query optimizer the most efficient way to execute the statement. Hints override the optimizer's execution plan, so you can use hints to improve performance by forcing the optimizer to use a more efficient plan. Hints are only supported for Oracle Database data sources.

Using the Administration Tool, you can add hints to a repository, in both online and offline modes, to optimize the performance of queries. When you add a hint to the repository, you associate it with Physical layer objects. When the object associated with the hint is queried, the Oracle BI Server inserts the hint into the SQL statement.

Table 7–6 shows the physical objects with which you can associate hints. It also shows the Administration Tool dialog that corresponds to the physical object. Each of these dialogs contains a **Hint** field, into which you can type a hint to add it to the repository.

Table 7–6 Physical Layer Objects That Accept Hints

Database Object	Dialog
Complex join	Complex Join
Physical foreign key	Physical Foreign Key
Physical table	Physical Table - General tab

Hints are only supported when the **Table Type** is set to **Physical Table**. For other table types, the hint text is ignored. For physical tables with a table type of **Select**, you can provide the hint text as part of the SQL statement entered in the **Default Initialization String** field.

How to Use Oracle Hints

This section provides a few examples of how to use Oracle hints with the Oracle BI Server. For more information about Oracle hints, see *Oracle Database SQL Language Reference* for the version of the Oracle Database that you use.

About the Index Hint

The Index hint instructs the optimizer to scan a specified index rather than a table. [Example 7-1](#) explains how you would use the Index hint.

Example 7-1 Index Hint

You find queries against the `ORDER_ITEMS` table to be slow. You review the execution plan of the query optimizer and find the `FAST_INDEX` index is not being used. You create an Index hint to force the optimizer to scan the `FAST_INDEX` index rather than the `ORDER_ITEMS` table. The syntax for the Index hint is `index(table_name, index_name)`. To add this hint to the repository, go to the Physical Table dialog in the Administration Tool and type the following text in the **Hint** field:

```
index(ORDER_ITEMS, FAST_INDEX)
```

About the Leading Hint

The Leading hint forces the optimizer to build the join order of a query with a specified table. The syntax for the Leading hint is `leading(table_name)`. If you were creating a foreign key join between the Products table and the Sales Fact table and wanted to force the optimizer to begin the join with the Products table, you would go to the Physical Foreign Key dialog in the Administration Tool and type the following text in the **Hint** field:

```
leading(Products)
```

About Performance Considerations for Hints

Hints that are well researched and planned can result in significantly better query performance. However, hints can also negatively affect performance if they result in a suboptimal execution plan. Follow these guidelines to create hints to optimize query performance:

- You should only add hints to a repository after you have tried to improve performance in the following ways:
 - Added physical indexes (or other physical changes) to the Oracle Database.
 - Made modeling changes within the server.
- Avoid creating hints for physical table and join objects that are queried often. If you drop or rename a physical object that is associated with a hint, you must also alter the hints accordingly.

Creating Hints

The following procedure explains how to add hints to the repository using the Administration Tool.

To create a hint:

1. In the Administration Tool, go to one of the following dialogs:
 - Physical Table—General tab

- Physical Foreign Key
 - Complex Join
2. Type the text of the hint in the **Hint** field and click **OK**.

For a description of available Oracle hints and hint syntax, see *Oracle Database SQL Language Reference* for the version of the Oracle Database that you use.

Note: Although hints are identified using SQL comment markers (`/*` or `--`), do not type SQL comment markers when you type the text of the hint. The Oracle BI Server inserts the comment markers when the hint is executed.

Displaying and Updating Row Counts for Physical Tables and Columns

When you request row counts, the Administration Tool retrieves the number of rows from the data source for all or selected tables and columns (distinct values are retrieved for columns) and stores those values in the repository. The time this process takes depends upon the number of row counts retrieved.

When updating all row counts, the Updating Row Counts window appears while row counts are retrieved and stored. If you click **Cancel**, the retrieve process stops after the in-process table (and its columns) have been retrieved. Row counts include all tables and columns for which values were retrieved before the cancel operation.

Updating all row counts for a large repository might take a long time to complete. Therefore, you sometimes might want to update only selected table and column counts.

Row counts are not available for the following:

- Stored Procedure object types
- XML data sources and XML Server data sources
- Multidimensional data sources
- Data sources that do not support the `COUNTDISTINCT` function, such as Microsoft Access and Microsoft Excel, or data sources for which `COUNT_STAR_SUPPORTED` has been disabled in the database features table
- In online mode, Update Row Count does not work with connection pools in which the session variables `:USER` and `:PASSWORD` are set as the user name and password.
In offline mode, the Set values for variables dialog appears so that you can populate the session variables `:USER` and `:PASSWORD`.
- In online mode, after importing or manually creating a physical table or column, the Oracle BI Server does not recognize the new objects until you check them in. Therefore, **Update Row Count** is not available in the menu until you check in these objects.

To display row counts in the Physical layer:

1. In the Administration Tool, select **Tools**, then select **Options**.
2. In the General tab of the Options dialog, select **Show row count in physical view**, and then click **OK**.

To update selected row counts in the Physical layer:

1. In the Physical layer of the Administration Tool, right-click a single table or column. You can select multiple objects and then right-click.
2. In the shortcut menu, select **Update Row Count**.

To update all row counts in the Physical layer:

1. In the Administration Tool, select **Tools**, then select **Update All Row Counts**. If the repository is open in online mode, the Check Out Objects window might open.
2. Click **Yes** to check out the objects.
Any row counts that have changed since the last update are refreshed.

Working with Logical Tables, Joins, and Columns

The Business Model and Mapping layer of the Oracle BI repository defines the business, or logical, model of the data and specifies the mapping between the business model and the Physical layer schemas. Business models are always dimensional, unlike objects in the Physical layer, which reflect the organization of the data sources. The Business Model and Mapping layer can contain one or more business models. Each business model contains logical tables, columns, and joins.

Even though similar terminology is used for logical table and physical table objects, such as the concept of keys, logical tables and joins in the Business Model and Mapping layer have their own set of rules that differ from those of relational models. For example, logical fact tables are not required to have keys, and logical joins can represent many possible physical joins.

Logical tables, joins, mappings, and other objects in the Business Model and Mapping layer are typically created automatically when you drag and drop objects from the Physical layer to a particular business model. After these objects have been created, you can perform tasks like creating additional logical joins, performing calculations and transformations on columns, and adding and removing keys from dimension and fact tables.

This chapter contains the following sections:

- [Creating the Business Model and Mapping Layer](#)
- [Working with the Business Model Diagram](#)
- [Creating and Managing Logical Tables](#)
- [Defining Logical Joins](#)
- [Creating and Managing Logical Columns](#)
- [Enabling Write Back On Columns](#)
- [Setting Up Display Folders in the Business Model and Mapping Layer](#)
- [Modeling Bridge Tables](#)

Creating the Business Model and Mapping Layer

After creating all of the elements of the Physical layer, you can drag tables or columns from the Physical layer to a business model in the Business Model and Mapping layer to create logical objects in the metadata.

This section contains the following topics:

- [Creating Business Models](#)
- [Automatically Creating Business Model Objects](#)
- [Duplicating a Business Model and Subject Area](#)

Creating Business Models

The Business Model and Mapping layer of the Administration Tool can contain one or more business models. A business model contains the business model definitions and the mappings from logical to physical tables for the business model.

When you work in a repository in offline mode, remember to save your repository from time to time. You can save a repository in offline mode even though the business models may be inconsistent.

To create a business model:

1. In the Administration Tool, right-click in the Business Model and Mapping layer below any existing objects.
2. Select the option **New Business Model** from the shortcut menu.
3. Specify a name for the business model.
4. New business models are disabled by default. If you want to make the corresponding Presentation layer available for queries, deselect **Disabled**.

Note: The business model should be consistent before you deselect this option.

5. Optionally, type a description of the business model.
6. Click **OK**.

After you create a business model, you can create business model objects by dragging and dropping objects from the Physical layer. See the next section for more information.

Automatically Creating Business Model Objects

To automatically map objects in the Business Model and Mapping layer to sources in the Physical layer, you can drag and drop Physical layer objects to a particular business model in the logical layer. When you drag a physical table to the Business Model and Mapping layer, a corresponding logical table is created. For each physical column in the table, a corresponding logical column is created. If you drag multiple tables at once, a logical join is created for each physical join, but only the first time the tables are dragged onto a new business model.

Automatically Creating Business Model Objects for Multidimensional Data Sources

Setting up objects in the Business Model and Mapping layer for multidimensional data sources is similar to setting up logical layer objects for a relational data source. To create the business model layer, you can drag and drop the Physical layer cube to the logical layer. Oracle Business Intelligence automatically creates a fully configured and consistent business model that retains metrics, attributes and dimensions.

Note: For Essbase data sources, it is recommended that you create a separate business model for each Essbase cube. To do this, drag each cube individually to the Business Model and Mapping layer.

Duplicating a Business Model and Subject Area

This feature lets you select a business model and its corresponding subject area (or a subject area and its corresponding business model), make a copy, and assign new names to the duplicates. Note that aliases are not copied.

To copy a business model and subject area:

1. Perform one of the following steps:
 - In the Business Model and Mapping layer of the Administration Tool, right-click a business model and select **Duplicate with Subject Area**.
 - In the Presentation layer of the Administration Tool, right-click a subject area and select **Duplicate with Business Model**.
2. In the Copy Business Model and Subject Area dialog, select the business model and corresponding subject area you want to copy.
3. Specify new names for the business model and subject area in the appropriate name fields, and then click **OK**.

The copied business model appears in the Business Model and Mapping layer, and the copied subject area appears in the Presentation layer.

Working with the Business Model Diagram


In addition to working with Business Model and Mapping layer objects in the middle pane of the Administration Tool, you can open the Business Model Diagram to see a graphical model of logical tables and joins.

To access the Business Model Diagram, right-click an object in the Business Model and Mapping layer (such as a dimension or fact table) and select **Business Model Diagram**. Then, select one of the following options:

- **Whole Diagram.** Displays all logical tables and joins in the business model.
- **Selected Tables Only.** Displays only the selected logical tables. Logical joins appear only if they exist between the objects that you selected. This option is only available when you select one or more logical tables.
- **Selected Tables and Direct Joins.** Displays the selected logical tables and any logical tables that join to the tables that you selected. This option is only available when you select one or more logical tables.
- **Selected Fact Tables and Dimensions.** Displays the selected logical tables and their associated logical dimensions. This option is only available when your selection includes at least one fact table.

Note that the Business Model Diagram displays only logical tables and joins. It does not display other Business Model and Mapping layer objects, such as business models, dimensions, or hierarchies. Joins are represented by a line with an arrow at the "one" end of the join.

Figure 8–1 Business Model Diagram



To add additional tables to the Business Model Diagram, leave the Business Model Diagram window open and then right-click the table or tables you want to add. Then, select **Business Model Diagram** and choose one of the display options.

Additional options are available in the right-click menu for the graphical tables and joins displayed in the Business Model Diagram. For example, you can delete objects or view their properties, or you can add additional related objects using the right-click options **Add Direct Joins**, **Add Tables Joined to Whole Selection**, and **Add All Joins**. You can also select **Find in Tree View** to locate a particular object in the Business Model and Mapping layer view in the middle pane, or check out objects in online mode.

You can also right-click an object in the Business Model Diagram view and select **Hide** to hide particular objects in the diagram. Note that this effect is temporary and does not persist.

See also the following sections:

- ["Using the Physical and Business Model Diagrams"](#) for information about zooming, panning, and controlling the layout of the tables
- ["Defining Logical Joins with the Business Model Diagram"](#) for information about defining logical joins

Creating and Managing Logical Tables

Logical tables exist in the Business Model and Mapping layer. The logical schema defined in each business model must contain at least two logical tables, and you must define relationships between them.

Each logical table has one or more logical columns and one or more logical table sources associated with it. You can change the logical table name, reorder the logical table sources, and configure the logical keys, both primary and foreign.

This section contains the following topics:

- [Creating Logical Tables](#)
- [Specifying a Primary Key in a Logical Table](#)
- [Reviewing Foreign Keys for a Logical Table](#)

Creating Logical Tables

Typically, you create logical tables by dragging and dropping a physical table from the Physical layer to a business model in the Business Model and Mapping layer. If a table does not exist in your physical schema, you need to create the logical table manually.

Drag and drop operations are usually the fastest method for creating objects in the Business Model and Mapping layer. If you drag and drop physical tables from the Physical layer to the Business Model and Mapping layer, the columns belonging to the table are also copied. After you drag and drop objects into the Business Model and Mapping layer, you can modify them in any way necessary without affecting the objects in the Physical layer.

When you drag physical tables (with key and foreign key relationships defined) to a business model, logical keys and joins are created that mirror the keys and joins in the Physical layer. This occurs only if the tables that you drag include the table with the foreign keys. Additionally, if you create new tables or subsequently drag additional tables from the Physical layer to the Business Model and Mapping layer, the logical mappings between the new or newly dragged tables and the previously dragged tables must be created manually.

See "[Defining Logical Joins with the Joins Manager](#)" and "[Defining Logical Joins with the Business Model Diagram](#)" for more information about joins.

To create a logical table by dragging and dropping:

1. In the Administration Tool, select one or more table objects in the Physical layer.
You must include the table with the foreign keys if you want to preserve the keys and joins from the Physical layer.
2. Drag and drop the table objects to a business model in the Business Model and Mapping layer.
When you drop them, the table objects, including the physical source mappings, are created automatically in the Business Model and Mapping layer.

To create a logical table manually:

1. In the Business Model and Mapping layer of the Administration Tool, right-click the business model in which you want to create the table and select **New Object > Logical Table**.
The Logical Table dialog appears.
2. In the General tab, type a name for the logical table.
3. If this is a lookup table, select the option **Lookup table**. A lookup table stores multilingual data corresponding to rows in the base tables. See "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization and lookup tables.

4. Optionally, type a description of the table.
5. Click **OK**.

After creating a logical table manually, you must create all keys and joins manually.

Creating and Managing Logical Table Sources

You can add a new logical table source, edit or delete an existing table source, create or change mappings to the table source, and define when to use logical tables sources and how content is aggregated. See [Chapter 10, "Managing Logical Table Sources \(Mappings\)"](#) for instructions about how to perform these tasks.

Specifying a Primary Key in a Logical Table

After creating tables in the Business Model and Mapping layer, you specify a primary key for each dimension table. Logical dimension tables must have a logical primary key. Logical keys can be composed of one or more logical columns.

Note: It is recommended that you do not specify logical keys for logical fact tables.

To specify a primary key in a logical table:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a table.
2. In the Logical Table dialog, select the Keys tab and then click **New**.
3. In the Logical Key dialog, type a name for the key and select the column that defines the key of the logical table.
4. Click **OK**.

Reviewing Foreign Keys for a Logical Table

It is recommended that you do not use foreign key joins in logical tables. If you must create these joins, you must first enable the option **Allow logical foreign key join creation** in the Options dialog. See ["Creating Logical Foreign Key Joins with the Joins Manager"](#) for more information.

The Foreign Keys tab of the Logical Table dialog exists so that you can view logical foreign keys you might have had in a previous release of Oracle Business Intelligence.

Defining Logical Joins

Relationships between logical tables are expressed by logical joins. Logical joins are *conceptual*, rather than physical, joins. In other words, they do not join to particular keys or columns. A single logical join can correspond to many possible physical joins.

A key property of a logical join is cardinality. Cardinality expresses how rows in one table are related to rows in the table to which it is joined. A one-to-many cardinality means that for every row in the first logical dimension table, there are 0, 1, or many rows in the second logical table. The Administration Tool considers a table to be a logical fact table if it is at the Many end of all logical joins that connect it to other logical tables.

Specifying the logical table joins is required so that the Oracle BI Server can have the necessary metadata to translate a logical request against the business model to SQL

queries against the physical data sources. The logical join information provides the Oracle BI Server with the many-to-one relationships between the logical tables. This logical join information is used when the Oracle BI Server generates queries against the underlying databases.

You do not need to create logical joins in the Business Model and Mapping layer if both of the following statements are true:

- You create the logical tables by simultaneously dragging and dropping all required physical tables to the Business Model and Mapping layer.
- The logical joins are the same as the joins in the Physical layer.

However, you will probably have to create some logical joins in the Business Model and Mapping layer, because you will rarely drag and drop all physical tables simultaneously except in very simple models.

You can create logical joins using either the Joins Manager or the Business Model Diagram. When you create a complex join in the Physical layer, you can specify expressions and the specific columns on which to create the join. When you create a logical join in the Business Model and Mapping layer, you cannot specify expressions or columns on which to create the join. The existence of a join in the Physical layer does not require a matching join in the Business Model and Mapping layer.

Note: It is recommended that you do not have foreign keys for logical tables. However, for backward compatibility, you can create logical foreign key joins using the Joins Manager if you select **Allow logical foreign key join creation** in the Options dialog.

A logical key for a fact table must be made up of the key columns that join to the attribute tables. Logical foreign key joins may be needed if the Oracle BI Server is to be used as an ODBC data source for certain third-party query and reporting tools.

This section contains the following topics:

- [Defining Logical Joins with the Business Model Diagram](#)
- [Defining Logical Joins with the Joins Manager](#)
- [Specifying a Driving Table](#)
- [Identifying Physical Tables That Map to Logical Objects](#)

Defining Logical Joins with the Business Model Diagram

The Business Model Diagram shows logical tables and any defined joins between them. You can use the Business Model Diagram to define logical joins between tables.

To define a logical join with the Business Model Diagram:

1. In the Administration Tool, right-click a business model and select **Business Model Diagram**, then select **Whole Diagram**.
2. Click the **New Join** button on the Administration Tool toolbar:



3. In the Business Model Diagram, left-click the first table in the join (the table representing many in the one-to-many join) to select it.

4. Move the cursor to the table to which you want to join (the table representing one in the one-to-many join), and then left-click the second table to select it.

The Logical Join dialog appears.

5. (Optional) To specify a driving table for the key, select a table from the **Driving table** list, and an applicable cardinality.


This option is useful for optimizing the manner in which the Oracle BI Server processes multi-database inner joins when one table is very small and the other table is very large. Do not select a driving table unless multi-database joins are going to occur. See "[Specifying a Driving Table](#)" for more information about driving tables.

Caution: Use extreme caution in deciding whether to specify a driving table. Driving tables are used for query optimization only under rare circumstances and when the driving table is extremely small (fewer than 1000 rows). Choosing a driving table incorrectly can lead to severe performance degradation.

6. Select the join type from the **Type** list, or keep the default value.
7. Set the **Cardinality** for each side of the join, or keep the default values.
8. Click **OK** to save your work.

In the Business Model Diagram, the join is represented by a line between the two selected tables, with an arrow at the "one" end of the join. [Figure 8–2](#) shows a join in the Business Model Diagram.

Figure 8–2 Join in the Business Model Diagram



Defining Logical Joins with the Joins Manager

You can use the Joins Manager to view logical join relationships and to create logical joins. You can also use the Joins Manager to create logical foreign key joins if you select **Allow logical foreign key join creation** in the Options dialog, although this is not recommended.

This section contains the following topics:

- [Creating Logical Joins with the Joins Manager](#)
- [Creating Logical Foreign Key Joins with the Joins Manager](#)

Creating Logical Joins with the Joins Manager

Logical joins are recommended over logical foreign key joins in the Business Model and Mapping layer.

To create a logical join with the Joins Manager:

1. In the Administration Tool, select **Manage**, then select **Joins**.
The Joins Manager dialog appears.
2. Select **Action > New > Logical Join**.

The Logical Join dialog appears.

3. Type a name for the logical join.
4. In the **Table** lists on the left and right side of the dialog, select the tables that the logical join references.
5. (Optional) To specify a driving table for the key, select a table from the **Driving** list, and an applicable cardinality.

This option is useful for optimizing the manner in which the Oracle BI Server processes multi-database inner joins when one table is very small and the other table is very large. Do not select a driving table unless multi-database joins are going to occur. See "[Specifying a Driving Table](#)" for more information about driving tables.

Caution: Use extreme caution in deciding whether to specify a driving table. Driving tables are used for query optimization only under rare circumstances and when the driving table is extremely small, that is, less than 1000 rows. Choosing a driving table incorrectly can lead to severe performance degradation.

6. Select the join type from the **Type** list, or keep the default value.
7. Set the **Cardinality** for each side of the join, or keep the default values.
8. Click **OK**.

Creating Logical Foreign Key Joins with the Joins Manager

Logical foreign key joins might be needed if the Oracle BI Server is to be used as an ODBC data source for certain third-party query and reporting tools. Typically, you should not create logical foreign keys. This capability is in the Administration Tool to provide compatibility with previous releases.

To create a logical foreign key join with the Joins Manager:

1. In the Administration Tool, select **Tools**, then select **Options**.
2. In the General tab of the Options dialog, select **Allow logical foreign key join creation**.
3. Click **OK**.
4. Select **Manage**, then select **Joins** to display the Joins Manager.
5. Select **Action > New > Logical Foreign Key**.
6. In the Browse dialog, double-click a table to display the Logical Foreign Key dialog.
7. Type a name for the foreign key.
8. In the **Table** list on the left side of the dialog, select the table that the foreign key references.
9. Select the columns in the left table that the foreign key references.
10. Select the columns in the right table that make up the foreign key columns.
11. (Optional) To specify a driving table for the key, select a table from the **Driving** list, and an applicable cardinality.

This option is useful for optimizing the manner in which the Oracle BI Server processes multi-database inner joins when one table is very small and the other table is very large. Do not select a driving table unless multi-database joins are going to occur. See ["Specifying a Driving Table"](#) for more information about driving tables.

Caution: Use extreme caution in deciding whether to specify a driving table. Driving tables are used for query optimization only under rare circumstances and when the driving table is extremely small, that is, less than 1000 rows. Choosing a driving table incorrectly can lead to severe performance degradation.

12. Select the join type from the **Type** list, or keep the default value.
13. Set the **Cardinality** for each side of the join, or keep the default values.
14. Enter an expression for the join, or click the **Expression Builder** button to define the expression in Expression Builder.
15. Click **OK** to save your work.

Specifying a Driving Table

You can specify a driving table for logical joins from the Logical Joins window. Driving tables are useful for optimizing the manner in which the Oracle BI Server processes cross-database joins when one table is very small and the other table is very large. Specifying driving tables leads to query optimization only when the number of rows being selected from the driving table is much smaller than the number of rows in the table to which it is being joined.

Caution: To avoid problems, only specify driving tables when the driving table is extremely small - less than 1000 rows.

When you specify a driving table, the Oracle BI Server uses it if the query plan determines that its use will optimize query processing. The small table (the driving table) is scanned, and parameterized queries are issued to the large table to select matching rows. The other tables, including other driving tables, are then joined together.

Caution: If large numbers of rows are being selected from the driving table, specifying a driving table could lead to significant performance degradation or, if the `MAX_QUERIES_PER_DRIVE_JOIN` limit is exceeded, the query terminates.

In general, driving tables can be used with inner joins, and for outer joins when the driving table is the left table for a left outer join, or the right table for a right outer join. Driving tables are not used for full outer joins. See ["Defining Logical Joins"](#) for instructions on specifying a driving table.

There are two entries in the database features table that control and tune driving table performance.

- `MAX_PARAMETERS_PER_DRIVE_JOIN`

This is a performance tuning parameter. In general, the larger its value, the fewer parameterized queries need to be generated. Values that are too large can result in parameterized queries that fail due to back-end database limitations. Setting the value to 0 (zero) turns off drive table joins.

- `MAX_QUERIES_PER_DRIVE_JOIN`

This is used to prevent runaway drive table joins. If the number of parameterized queries exceeds its value, the query is terminated and an error message is returned to the user.

Identifying Physical Tables That Map to Logical Objects

The Physical Diagram shows the physical tables that map to the selected logical object and the physical joins between each table.

One of the joins options, **Object(s) and Direct Joins within Business Model**, is unique to the logical layer. It creates a physical diagram of the tables that meet both of the following conditions:

- Tables in the selected objects and tables that join directly
- Tables that are mapped (exist in logical table sources in the business model) in the business model

To open the Physical Diagram for a logical object:

1. In the Business Model and Mapping layer of the Administration Tool, right-click a business model, logical table, or logical table source.
2. Select **Physical Diagram** and then one of the joins options.
3. Click and drag any object to more clearly view the relationship lines, such as one-to-many.

Creating and Managing Logical Columns

Many logical columns are automatically created by dragging tables from the Physical layer to the Business Model and Mapping layer. Other logical columns, especially ones that involve calculations based on other logical columns, can be created later.

Logical columns are displayed in a tree structure expanded out from the logical table to which they belong. If the column is a primary key column or participates in a primary key, the column is displayed with a key icon. If the column has an aggregation rule, it is displayed with a ruler icon. You can also reorder logical columns in the Business Model and Mapping layer.

This section contains the following topics:

- [Creating Logical Columns](#)
- [Basing the Sort for a Logical Column on a Different Column](#)
- [Enabling Double Column Support by Assigning a Descriptor ID Column](#)
- [Creating Derived Columns](#)
- [Setting Default Levels of Aggregation for Measure Columns](#)
- [Associating an Attribute with a Logical Level in Dimension Tables](#)
- [Moving or Copying Logical Columns](#)

Creating Logical Columns

The following procedure explains how to create logical columns in the Business Model and Mapping layer.

To create a logical column:

1. In the Business Model and Mapping layer, right-click a logical table.
2. From the shortcut menu, select **New Object**, then select **Logical Column**.
3. In the General tab, type a name for the logical column.
The name of the business model and the associated logical table appear in the **Belongs to Table** field.
4. Select **Writeable** to enable write back for this column. See ["Enabling Write Back On Columns"](#) for more information.
5. Optionally, you can assign a different column on which to base the sort order for a column. See ["Basing the Sort for a Logical Column on a Different Column"](#) for details.
6. Optionally, you can assign a descriptor ID column for this column. See ["Enabling Double Column Support by Assigning a Descriptor ID Column"](#) for details.
7. Optionally, on the Column Source tab, you can specify that this logical column is derived from other logical columns. See ["Creating Derived Columns"](#) for details.
8. Optionally, on the Aggregation tab, you can set column aggregation. See ["Setting Default Levels of Aggregation for Measure Columns"](#) for details.
9. Optionally, on the Levels tab, you can associate attributes with a logical level. Measures can be associated with levels from multiple dimensions and always aggregate to the levels specified. See ["Associating an Attribute with a Logical Level in Dimension Tables"](#) for details.
10. Click **OK**.

Basing the Sort for a Logical Column on a Different Column

For a logical column, you can specify a different column on which to base the sort. This changes the sort order of a column when you do not want to order the values lexicographically. Lexicographical sort arranges the results in alphabetic order such as in a dictionary. In this type of sort, numbers are ordered by their alphabetic spelling and not divided into a separate group.

For example, if you sorted on month (using a column such as `MONTH_NAME`), the results would be returned as February, January, March, and so on, in lexicographical sort order. However, you might want months to be sorted in chronological order. Therefore, your table should have a month key (such as `MONTH_KEY`) with values of 1 (January), 2 (February), 3 (March), and so on. To achieve the desired sort, you set the Sort order column field for the `MONTH_NAME` column to be `MONTH_KEY`. Then, a request to order by `MONTH_NAME` would return January, February, March, and so on.

To assign a different column on which to base the sort order for a column:

1. In the Logical Column dialog, in the General tab, click **Set** next to the **Sort order column** field.
2. In the Browse dialog, select a column.
3. To view the column details, click **View** to open the Logical Column dialog for that column, and then click **Cancel**.

You can make some changes in this dialog. If you make changes, click **OK** to accept the changes instead of **Cancel**.

4. In the Browse dialog, click **OK**.

Enabling Double Column Support by Assigning a Descriptor ID Column

When multilingual columns are based on a lookup function, it is common to specify the non-translated lookup key column as the descriptor ID column of the translated column. Assigning a descriptor ID column enables Double Column Support, a feature which helps in defining language-independent filters. For example, in Answers, users see the display column, but the query filters on the hidden descriptor ID column.

For more information, see "Supporting Multilingual Data" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Note that double columns are also used for other purposes, like modeling spatial columns.

To assign a Descriptor ID column to a display column:

1. In the Logical Column dialog, in the General tab, click **Set** next to the **Descriptor ID column** field.
2. In the Browse dialog, select a key column.
3. To view the column details, click **View** to open the Logical Column dialog for that column, and then click **Cancel**.

You can make some changes in this dialog. If you make changes, click **OK** to accept the changes instead of **Cancel**.

4. In the Browse dialog, click **OK**.

Creating Derived Columns

Some columns are derived from other logical columns as a way to apply post-aggregation calculations to measures. To do this, you specify the derived column expression in the Column Source tab of the Logical Column dialog.

You can also create a set of derived columns using the Calculation Wizard. See "[Using the Calculation Wizard](#)" for more information.

Note that if the parameter `PREVENT_DIVIDE_BY_ZERO` is set to `YES` in `NQSCONFIG.INI`, the Oracle BI Server prevents errors in divide-by-zero situations, even for Answers column calculations. The Oracle BI Server creates a divide-by-zero prevention expression using `nullif()` or a similar function when it writes the physical SQL. Because of this, you do not have to use `CASE` statements to avoid divide-by-zero errors, as long as `PREVENT_DIVIDE_BY_ZERO` is set to `YES` (the default value).

See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about `NQSCONFIG.INI` settings.

You can also apply calculations pre-aggregation. See "[Defining Physical to Logical Table Source Mappings and Creating Calculated Items](#)" for more information.

To specify a derived column:

1. In the Logical Column dialog, select the Column Source tab.
2. Select the option **Derived from existing columns using an expression**.
3. Click the **Expression Builder** button to open Expression Builder.

4. In the Expression Builder - Derived logical column dialog, specify the expression from which the logical column should be derived.

Note: To optimize performance, do not define aggregations in Expression Builder. Instead, use the Aggregation tab of the Logical Column dialog. See ["Setting Default Levels of Aggregation for Measure Columns"](#) for more information.

5. Click **OK**.

Note that you can display data from multilingual database schemas by using Expression Builder to create a lookup function. For more information, see "Supporting Multilingual Data" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Configuring Logical Columns for Multicurrency Support

You can configure logical columns so that Oracle Business Intelligence users can select the currency in which they prefer to view currency columns in analyses and dashboards. You can set up this feature so that all users see the same static list of currency options, or you can provide a dynamic list of currency options that changes based on a Logical SQL statement you specify.

To configure logical columns for multicurrency support:

1. Create a session variable named `PREFERRED_CURRENCY`, along with an initialization block to use in the variable. Make sure to select **Enable any user to set the value** when you create the session variable. Note that when you use session variables in an expression for Oracle BI Presentation Services, you must preface their names with `NQ_SESSION`.

See ["Creating Session Variables"](#) and ["Creating Initialization Blocks"](#) for detailed information about setting up session variables and initialization blocks.

2. Edit any logical columns that display currency values to use the appropriate conversion factor using the `PREFERRED_CURRENCY` session variable. To do this, double-click the appropriate logical column in the Business Model and Mapping layer, select the Column Source tab, and create a derived expression that uses the `PREFERRED_CURRENCY` variable.

For example, the following logical column expression uses the value of the `NQ_SESSION.PREFERRED_CURRENCY` variable to switch between different currency columns. Note that the currency columns are expected to have the appropriate converted values.

```
INDEXCOL( CASE VALUEOF(NQ_SESSION.PREFERRED_CURRENCY) WHEN 'gc1' THEN 0
WHEN 'gc2' THEN 1 WHEN 'orgc' THEN 2 WHEN 'lc1' THEN 3 ELSE 4 END,
"Paint"."Sales Facts"."USDCurrency",
"Paint"."Sales Facts"."DEMCurrency" ,
"Paint"."Sales Facts"."EuroCurrency" ,
"Paint"."Sales Facts"."JapCurrency" ,
"Paint"."Sales Facts"."USDCurrency" )
```

3. If you want to provide a dynamic list of currency options, create a table in your data source that provides the entries you want to display for the user-preferred currency. This table must include the following columns:

- The first column contains the values used to set the session variable `PREFERRED_CURRENCY`. Each value in this column is a string that uniquely identifies the currency (for example, `gc2`).
- The second column contains currency tags from the file `currencies.xml`. The `displayMessage` values for each tag are used to populate the Currency box and currency prompts (for example, `int:euro-1`). The `currencies.xml` file is located in `ORACLE_HOME\bifoundation\web\display`.
- You can optionally provide a third column that contains the values used to set the presentation variable `currency.userPreference`. Each value in this column is a string that identifies the currency (for example, `Global Currency 2`). If you omit this column, then the values for the `displayMessage` attributes for the corresponding currency tags in the `currencies.xml` file are used.

Table 8–1 shows a sample table with user-preferred currency entries.

Table 8–1 Sample Table for Dynamically Displaying the Preferred Currency

UserPreference	CurrencyTag	UserPreferenceName
char	char	char
orgc1	loc:en-BZ	Org currency
gc2	int:euro-1	Global currency 2
lc1	int:DEM	Ledger currency
gc1	int:USD	Global Currency 1

Additional configuration is required in Oracle BI Presentation Services to enable this feature. For full information about the Oracle BI Presentation Services configuration, see "Defining User-Preferred Currency Options" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Setting Default Levels of Aggregation for Measure Columns

You need to specify aggregation rules for mapped logical columns that are measures. Aggregation should only be performed on measure columns, with the possible exception of the aggregation `COUNT` and `COUNTDISTINCT`. Measure columns should exist only in logical fact tables.

You can optionally select different aggregation rules for different dimensions that are associated with this logical column. For example, if someone queries the aggregate column along with one dimension, you may want to use one type of aggregation rule, whereas with another dimension, you may want to use a different aggregation rule.

When the default aggregation rule is `Count Distinct`, you can optionally specify an override aggregation expression for specific logical table sources. For example, you may want to specify override aggregation expressions when you are querying different aggregate table sources that already contain some level of aggregation. If you do not specify any override, then the default rule prevails.

You can choose the aggregation rule `Evaluate_Aggr` to enable queries to call custom functions in the data source. For information about this function and other aggregation rules, see [Appendix C](#). See also "Defining Aggregation Rules for Multidimensional Data Sources" for additional information about setting aggregation for multidimensional sources.

By default, data is considered sparse. However, on rare occasions you might have a logical table source with dense data. A logical table source is considered to have dense

data if it has a row for every combination of its associated dimension levels. When setting up aggregate rules for a measure column, you can specify that data is dense only if *all* the logical table sources to which it is mapped are dense.

To specify a default aggregation rule for a measure column:

1. In the Business Model and Mapping layer, double-click a logical column.
2. In the Logical Column dialog, click the Aggregation tab.
3. In the Aggregation tab, choose one of the following options:
 - For measures in which the additivity is the same in all dimensions (in other words, for fully-additive or non-additive measures), select one of the aggregate functions from the **Default Aggregation Rule** list.

The function you select is always applied when a user or an application requests the column in a query, unless an override aggregation expression has been specified.

When you select Count Distinct as the default aggregation rule, you can specify an override aggregation expression for specific logical table sources. Choose this option when you have more than one logical table source mapped to a logical column and you want to apply a different aggregation rule to each source.

Click the **Add** button to select logical table sources for which you want to specify individual aggregation rules. In the Browse dialog, select the logical table source you want to add, and click **OK**. Then, in the **Formula** list for that logical table source, select the aggregation rule you want to use.

- Select **Based on dimensions** if your measure has different additivity for different dimensions (in other words, for semi-additive measures). For example, select this option for inventory units that sum in all dimensions except time. See "[Setting Up Dimension-Specific Aggregate Rules for Logical Columns](#)" for more information about this feature.

Click the **Add** button to select additional dimensions for which you want to specify aggregation rules. In the Browse dialog, select the dimension you want to add, and then click **OK**. Then, in the **Formula** list for that dimension, select the aggregation rule you want to use, or click the **Expression Builder** button to build the aggregation rule using Expression Builder.

The **Data is dense** option appears when you select **Based on dimensions**. Select this option only if *all* the logical table sources to which this column is mapped are dense.

Caution: Selecting **Data is dense** indicates that all sources to which this column is mapped have a row for every combination of dimension levels that they represent. Selecting this option when any table source that is used by this column does not contain dense data will return incorrect results.

4. Click **OK**.

Setting Up Dimension-Specific Aggregate Rules for Logical Columns

The majority of measures have the same aggregation rule for each dimension. However, some measures can have different aggregation rules for different dimensions. For example, bank balances might be averaged over time but summed

over the individual accounts. The Oracle BI Server lets you configure dimension-specific aggregation rules. You can specify one aggregation rule for a given dimension and specify other rules to apply to other dimensions.

You need to configure dimensions in the Business Model and Mapping layer to set up dimension-specific aggregation. For more information about setting up aggregate navigation, see [Chapter 10](#).

To specify dimension-specific aggregation rules for a single logical column:

1. In the Business Model and Mapping layer, double-click a logical column.
2. In the Logical Column dialog, click the Aggregation tab.
3. In the Aggregation tab, select **Based on dimensions**.
4. In the Browse dialog, select a dimension over which you want to aggregate, and then click **OK**.
5. In the Aggregation tab, from the **Formula** list, select a rule.
After selecting rules for specified dimensions, set the aggregation rule for any remaining dimensions by using the dimension labeled **Other**.
6. If you need to create more complex formulas, click the **Expression Builder** button to the right of the **Formula** column to open Expression Builder.
7. If you have multiple dimensions, you can click **Up** or **Down** to change the order in which the dimension-specific rules are performed.

When calculating the measure, aggregation rules are applied in the order (top to bottom) established in the dialog.

8. Click **OK**.

To specify dimension-specific aggregation rules for multiple logical fact columns:

1. In the Business Model and Mapping layer, select multiple logical fact columns.
2. Right-click and select **Set Aggregation**.
You must select more than one column to see the **Set Aggregation** menu item. Also note that **Set Aggregation** does not appear if one or more of the columns you select is a derived column.
3. In the Aggregation dialog, select or clear **All columns the same**.
This option is selected by default. When selected, you can set aggregation rules that apply to all selected columns. If you clear this option, you can set aggregation rules separately for each selected column.
4. In the Aggregation tab, select **Based on dimensions**.
5. In the Browse dialog, select a dimension over which you want to perform aggregation, and then click **OK**.
After setting up the rule for a dimension, specify aggregation rules for any other dimensions in the entry labeled **Other**.
6. Click the **Expression Builder** button to the right of the **Formula** column.
7. In the Expression Builder - Aggregate dialog, from the **Formula** list, select the aggregation to perform over the dimension.
8. To change the order in which the dimension-specific rules are performed, click **Up** or **Down**, and then click **OK**.

When calculating the measure, aggregation rules are applied in the order (top to bottom) established in the dialog.

Defining Aggregation Rules for Multidimensional Data Sources

This section describes best practices for defining aggregation rules for logical measures sourced from Essbase, Oracle OLAP, and other multidimensional data sources, like Microsoft Analysis Services and SAP/BW.

By default, when you import Essbase and some other multidimensional cubes into the Physical layer, Oracle Business Intelligence cannot read the aggregation rules set within the data source. Because of this, the measures are imported automatically with the default aggregation rule of **External Aggregation**. Note that this rule is only available for multidimensional data sources.

External Aggregation means that the Oracle BI Server is not aware of the underlying aggregation rule for the specific measure and will not compute it internally. Instead, the Oracle BI Server will always ship the query to the underlying multidimensional data source for aggregation.

Because the underlying data sources are extremely efficient, pushing the aggregation rules down to the data source ensures that the Oracle BI Server returns the results without adding any additional overhead in processing. However, it is recommended that you update the aggregation rule for each measure in Oracle Business Intelligence with the corresponding aggregation rule defined in the data source. Doing so ensures that the Oracle BI Server can do additional computations when needed. There is no query performance impact, since the Oracle BI Server still pushes down optimized queries wherever possible.

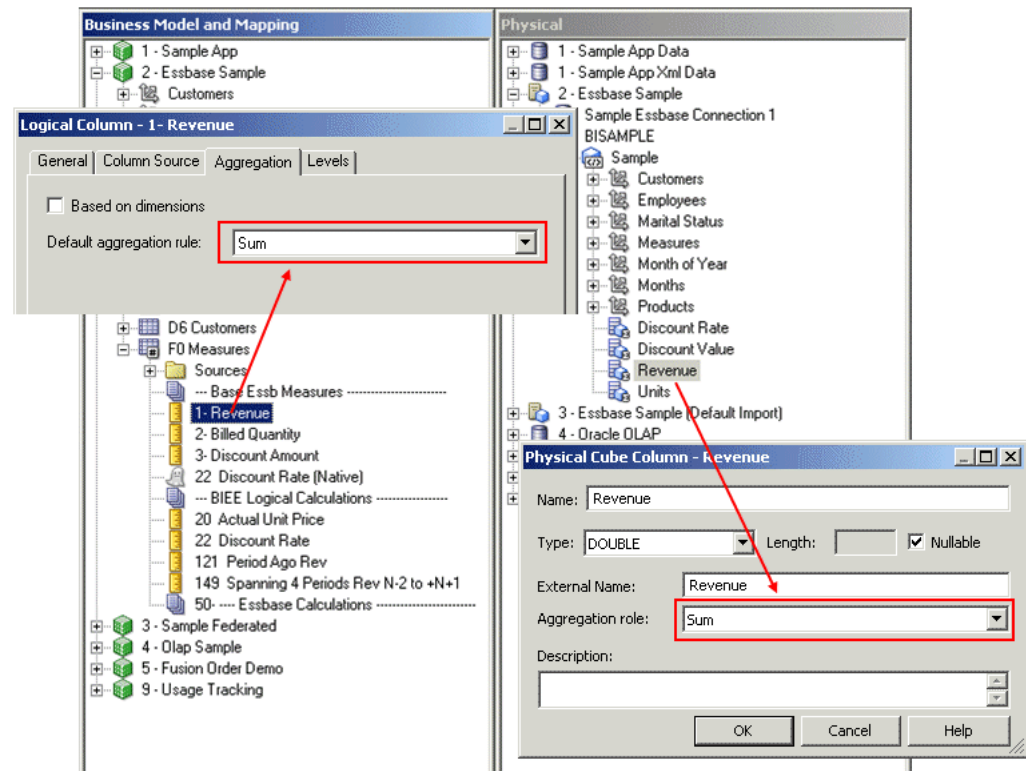
Note: If the Oracle BI Server needs to do additional aggregation for a particular query, and the aggregation rule is set to the default of External Aggregation, the server returns the following error:

An external aggregate is found in an outer query block.

This error occurs because the Oracle BI Server cannot read the aggregation rule in the underlying data source. To ensure that correct results are returned for these queries, you should change the aggregation rules set in the Oracle BI repository to match the aggregation rules set in the underlying data source.

You must ensure that the aggregation rule defined in Oracle Business Intelligence matches the rule in the underlying data source. Also, you must set the appropriate aggregation rule in both the Physical layer and Business Model and Mapping layer, as shown in [Figure 8-3](#).

Figure 8-3 Setting Aggregation Rules in the Physical and Business Model and Mapping Layers for Multidimensional Sources



For custom aggregations or aggregations which do not have a corresponding function within the Oracle BI Server, it is recommended to leave the aggregation as External Aggregation for both the physical measure column and its corresponding logical measure column.

Note: For Oracle OLAP data sources, you do not explicitly set Physical layer aggregation rules for Oracle OLAP columns. Because of this, you only need to set the aggregation rule for Oracle OLAP columns in the Business Model and Mapping layer.

In addition, if a query requests an aggregate that does not exist in the Oracle OLAP data source, and the aggregation rule is set to External Aggregation, then the Oracle BI Server returns an error. To avoid this error, make sure to explicitly set the aggregation rule for the Oracle OLAP column in the Business Model and Mapping layer.

If you do not explicitly set the aggregation rule for Oracle OLAP columns to something other than External Aggregation, requests from Oracle BI Presentation Services custom groups will fail, because custom groups always request aggregates that do not exist in the data source.

Associating an Attribute with a Logical Level in Dimension Tables

Attributes can be associated with a logical level by selecting the dimensional level on the Levels tab. Measures can be associated with levels from multiple dimensions and always aggregate to the levels specified.

Dimensions appear in the Dimensions list. If this attribute is associated with a logical level, the level appears in the Levels list.

Another way to associate a measure with a level in a dimension is to expand the dimension tree in the Business Model and Mapping layer, and then use drag-and-drop to drop the column on the target level. For more information about level-based measures, see [Example 9-1](#).

To associate a measure with a logical level in a dimension:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a logical column.
2. In the Logical Column dialog, click the Levels tab.
3. In the Levels tab, click the **Logical Level** field for the dimension from which you want to select a logical level.

In the Levels tab, in the levels list, you can sort the rows (toggle between ascending order and descending order) by clicking a column heading.

4. In the **Logical Level** list, select the level.
5. Repeat this process to associate this measure with other logical levels in other dimensions.

To remove the association between a dimension and a measure:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a logical column.
2. In the Logical Column dialog, click the Levels tab.
3. In the Levels tab, select the row for the association you want to remove and click **Delete**.
4. Click **OK**.

Moving or Copying Logical Columns

By default, dragging and dropping a logical column from one table to another moves the logical column. If a column with the same name already exists, the new column is renamed (for example, mycolumn#1).

You can also choose the option **Prompt when moving logical columns** in the Options dialog to cause the Sources for moved columns dialog to be displayed when you drag and drop a logical column. This dialog gives you options about the drag and drop behavior.

See "[Setting Preferences](#)" for more information about selecting the **Prompt when moving logical columns** option.

To move or copy logical columns using the Sources for moved columns dialog:

1. In the Business Model and Mapping layer, drag and drop a logical column to a different logical table. You can select multiple columns to move.
2. In the Sources for moved columns dialog, in the **Action** area, select an action.
3. If you select **Ignore**, no logical source is added in the Sources folder of the destination table.
4. If you select **Create new**, a copy of the logical source associated with the logical column is created in the Sources folder of the destination table.

5. If you select **Use existing**, in the **Use existing** list, you must select a logical source from the Sources folder of the destination table.

The column that you moved or copied is associated with this logical source.

Enabling Write Back On Columns

You can configure individual logical columns so that users in Oracle BI Presentation Services can update column data and write the changes back to the data source. To enable write back on a particular column, you must select the **Writeable** option for the logical column, and enable the **Read/Write** permission for the corresponding presentation column. You must also disable caching on the corresponding physical table.

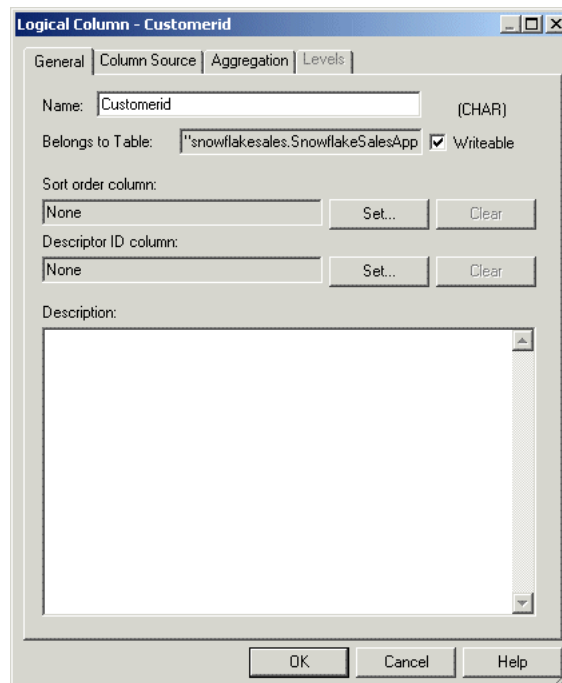
Additional tasks to enable write back need to be performed in Oracle BI Presentation Services. See "Configuring for Write Back in Analyses and Dashboards" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for full information.

To enable write back for a particular column:

1. In the Administration Tool, in the Physical layer, double-click the physical table that contains the column for which you want to enable write back.
2. On the General tab of the Physical Table dialog, ensure that **Cacheable** is not selected. Deselecting this option ensures that Oracle BI Presentation Services users can see updates immediately.
3. In the Business Model and Mapping layer, double-click the corresponding logical column. The Logical Column dialog opens.

Figure 8–4 shows the Logical Column dialog.

Figure 8–4 Logical Column Dialog with Writeable Option Selected




4. Select **Writeable**, then click **OK**.

5. In the Presentation layer, double-click the column that corresponds to the logical column for which you enabled write back. The Presentation Column dialog opens.
6. Click **Permissions**.
7. Select the **Read/Write** permission for the appropriate users and application roles.

Figure 8–5 shows the Permissions dialog.

Figure 8–5 Permissions Dialog for Presentation Layer Column with Read/Write Option Selected



8. Click **OK** in the Permissions dialog.
9. Click **OK** in the Presentation Column dialog.

Setting Up Display Folders in the Business Model and Mapping Layer

You can create display folders to organize objects in the Business Model and Mapping layer. They have no effect on query processing. After you create a display folder, the selected tables and dimensions appear in the folder as a shortcut and in the business model tree as the object. You can hide the objects so that you only view the shortcuts in the display folder. See the information about the Repository tab of the Options dialog in "[Setting Preferences](#)" for more information about hiding these objects.

Note: Deleting a table in a display folder deletes only the shortcut to that object. When you delete a column in a display folder, however, the column is actually deleted.

To set up a logical display folder:

1. In the Business Model and Mapping layer of the Administration Tool, right-click a business model and select **New Object**, then select **Logical Display Folder**.
2. In the Logical Display Folder dialog, in the Tables tab, type a name for the folder.
3. To add tables to the display folder, click **Add**. In the Browse dialog, select the fact or dimension tables you want to add to the folder and click **Select**.

Alternatively, you can drag one or more logical tables to the display folder after you close the dialog.

4. To add dimensions to the display folder, click the Dimensions tab and click **Add**. In the Browse dialog, select the dimensions that you want to add to the folder and click **Select**.

Alternatively, you can drag one or more dimensions to the display folder after you close the dialog.

5. Click OK.

Modeling Bridge Tables

A bridge table enables you to resolve many-to-many relationships between tables. For example, you might hold information about employees in an Employees table, and information about the jobs they do in a Jobs table. However, an organization's employees can have multiple jobs, and the same job can be performed by multiple employees. This situation would result in a many-to-many relationship between the Employees table and the Jobs table.

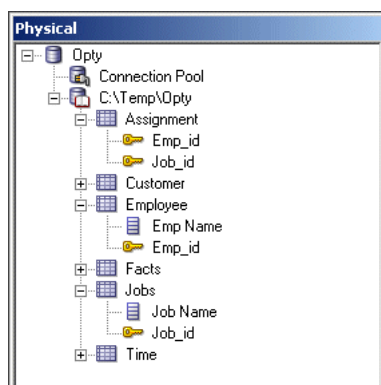
To resolve the many-to-many relationship, you can create a bridge table (or intermediate table) called Assignments. Each row in the Assignments table is unique, representing one employee doing one job. If an employee has several jobs, there are several rows in the Assignments table for that employee. If a job is done by several employees, there are several rows in the Assignments table for that job. The primary key of the Assignments table is a composite key, made up of a column containing the employee ID and a column containing the job ID.

By acting as a bridge table between the Job and Employee tables, the Assignments table enables you to resolve the many-to-many relationship between Employees and Jobs into:

- A one-to-many relationship between Employees and Assignments
- A one-to-many relationship between Assignments and Jobs

Figure 8–6 shows a Physical layer view of the example bridge and associated dimension tables described in the preceding paragraphs.

Figure 8–6 Example Bridge and Associated Tables in the Physical Layer



Note that "Weight Factor" should be included as an additional column in the bridge table and calculated during ETL for efficient query processing.

The following sections explain how to model bridge tables in the Physical and Business Model and Mapping layers:

- [Creating Joins in the Physical Layer for Bridge and Associated Dimension Tables](#)
- [Modeling the Associated Dimension Tables in a Single Dimension](#)
- [Modeling the Associated Dimension Tables in Separate Dimensions](#)

Creating Joins in the Physical Layer for Bridge and Associated Dimension Tables


To model bridge tables in the Physical layer, create joins between the bridge table and the associated dimension tables.

To create physical joins for a bridge table and its associated tables:

1. In the Administration Tool, in the Physical layer, select the fact, bridge, and associated dimension tables. Then, right-click the objects and select **Physical Diagram**, and then choose **Selected Object(s) Only**.
2. With the Physical Diagram displayed, click **New Join** on the toolbar. Then, select the bridge table, and then select one of the dimension tables.
3. Click **OK** in the Physical Foreign Key dialog.
4. Repeat steps 2 and 3 for the other associated dimension table.
5. Ensure that one of the associated dimension tables is joined to the fact table.

Figure 8–7 shows joins between the example Physical tables in the Physical Diagram.

Figure 8–7 Joins Between the Example Tables in the Physical Diagram



Modeling the Associated Dimension Tables in a Single Dimension

In the Business Model and Mapping layer, you can choose to model the two dimension tables associated with a bridge table in a single dimension, or in two separate dimensions. To model the associated dimension tables in one dimension, create a second logical table source that maps to the bridge table and the other dimension table, and then add columns from the other dimension table.

Providing two separate logical table sources makes queries more efficient, because it ensures that queries against a single dimension table do not involve the bridge table.

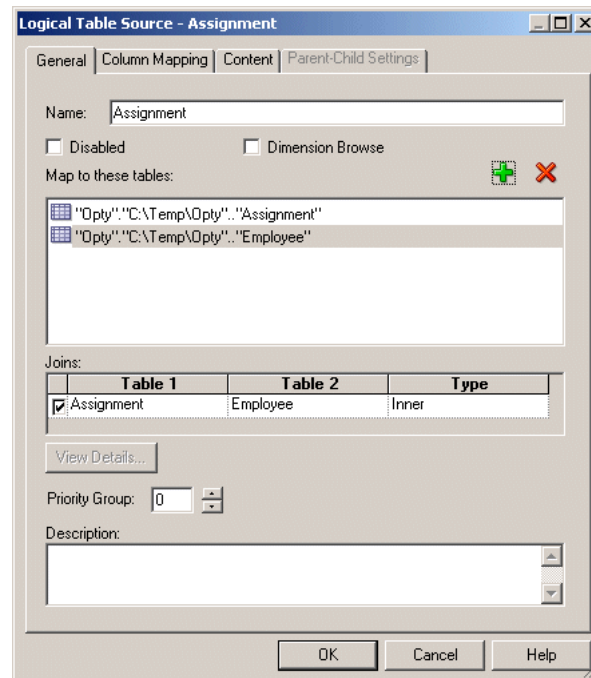
To model the dimension tables associated with a bridge table in a single dimension:

1. Drag objects from the Physical layer to the Business Model and Mapping layer, *except* the bridge table and the associated dimension table that is not joined to the fact table. For the example described in the previous sections, you would drag all objects except for the Assignment and Employee tables.
2. In the Business Model and Mapping layer, right-click the dimension table that is joined to the fact table (Jobs in our example) and select **New Object**, then select **Logical Table Source**.
3. In the Logical Table Source dialog, provide a name for the new bridge table source. It is a good practice to use the bridge table name as the name of the source (for example, Assignment).

4. Click the **Add** button in the upper right corner of the Logical Table Source dialog. Then, select the bridge table from the **Name** list (**Assignment** in our example) and then click **Select**.
5. Click the **Add** button again and select the associated dimension table that is not joined to the fact table (**Employee** in our example) and then click **Select**.
6. Click **OK** in the Logical Table Source dialog.

Figure 8–8 shows the Logical Table Source dialog for the bridge table source.

Figure 8–8 Logical Table Source Dialog for Bridge Table Source



7. Drag columns from the dimension table that is not joined to the fact table (Employees in our example) from the Physical layer to the logical table source that you just created.

You can now create dimensions based on your logical tables, including the logical table with the bridge table source.

Modeling the Associated Dimension Tables in Separate Dimensions

As an alternative to modeling the two dimension tables associated with a bridge table in a single dimension, you can choose to model them in separate dimensions. To do this, create a logical join between the fact table and the dimension table that is not physically joined to the fact table, and then modify the logical table source for that same dimension table to add the other table mappings.


To model the dimension tables associated with a bridge table in separate dimensions:

1. Drag objects from the Physical layer to the Business Model and Mapping layer. Because you want to model the dimension tables in separate dimensions, drag both of the dimension tables associated with the bridge table. You do not need to drag and drop the bridge table object.

2. In the Business Model and Mapping layer, select the fact table and the two dimension tables that are associated with the bridge table (Facts, Employee, and Jobs in our example). Then, right-click the objects and select **Business Model Diagram**, and then choose **Selected Tables Only**.
3. With the Business Model Diagram displayed, click **New Join** on the toolbar. Then, select the fact table, and then select the dimension table not currently joined to the fact table.
4. Click **OK** in the Logical Join dialog.

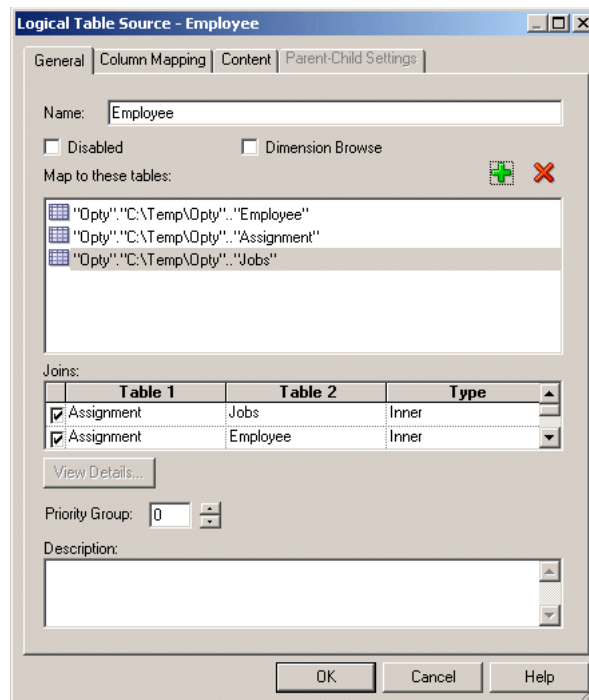
Figure 8–9 shows joins between the example logical tables in the Business Model Diagram.

Figure 8–9 Joins Between the Example Tables in the Business Model Diagram



5. Double-click the logical table source for the logical table for which you created the logical join (Employee in our example).
6. Click the **Add** button in the upper right corner of the Logical Table Source dialog. Then, select the bridge table from the **Name** list (**Assignment** in our example) and then click **Select**.
7. Click the **Add** button again and select the other associated dimension table (**Jobs** in our example) and then click **Select**.
8. Click **OK** in the Logical Table Source dialog.

Figure 8–10 shows the Logical Table Source dialog for the modified dimension table source.

Figure 8–10 Logical Table Source Dialog for Dimension Table Source

You can now create dimensions based on your logical tables, including both logical tables associated with the bridge table.

Working with Logical Dimensions

In the Business Model and Mapping layer, a dimension object represents a hierarchical organization of logical columns (attributes). One or more logical dimension tables can be associated with at most one dimension object. Common dimensions might be time periods, products, markets, customers, suppliers, promotion conditions, raw materials, manufacturing plants, transportation methods, media types, and time of day. Note that dimensions exist in the Business Model and Mapping (logical) layer and in the Presentation layer.

In each dimension, you organize logical columns into the structure of the hierarchy. This structure represents the organization rules and reporting needs required by your business. It also provide the metadata that the Oracle BI Server uses to drill into and across dimensions to get more detailed views of the data.

There are two types of logical dimensions: dimensions with level-based hierarchies (structure hierarchies), and dimensions with parent-child hierarchies (value hierarchies). Level-based hierarchies are those in which members are of several types, and members of the same type occur only at a single level. In parent-child hierarchies, members all have the same type. Oracle Business Intelligence also supports a special type of level-based dimension, called a time dimension, that provides special functionality for modeling time series data.

Because dimensions for multidimensional data sources are defined in the source, they do not require as much work compared with dimensions in other data sources. For example, you do not create dimension level keys. A dimension is specific to a particular multidimensional data source (it cannot be used by more than one) and cannot be created and manipulated individually. Additionally, each cube in the data source should have at least one dimension and one measure in the Business Model and Mapping layer.

You can expose logical dimensions to Oracle BI Answers users by creating presentation hierarchy objects that are based on particular logical dimensions. Creating hierarchies in the Presentation layer enables users to create hierarchy-based queries. See "[Working with Presentation Hierarchies and Levels](#)" for more information.

Note that you can also expose dimension hierarchies by adding one or more columns from each hierarchy level to a subject area in the Presentation layer. Oracle BI Answers supports drill-down on these hierarchical columns.

This chapter contains the following topics:

- [Creating and Managing Dimensions with Level-Based Hierarchies](#)
- [Creating and Managing Dimensions with Parent-Child Hierarchies](#)
- [Modeling Time Series Data](#)

Creating and Managing Dimensions with Level-Based Hierarchies

Each business model can have one or more dimensions, each dimension can have one or more logical levels, and each logical level has one or more attributes (columns) associated with it.

The following sections explain how to create dimensions:

- [About Level-Based Hierarchies](#)
- [Manually Creating Dimensions, Levels, and Keys with Level-Based Hierarchies](#)
- [Automatically Creating Dimensions with Level-Based Hierarchies](#)
- [Populating Logical Level Counts Automatically](#)

About Level-Based Hierarchies

A dimension contains two or more logical levels. The recommended sequence for creating logical levels is to create a Grand Total level and then create child levels, working down to the lowest level. The following are the parts of a dimension:

- **Grand Total level.** A special level representing the grand total for a dimension. Each dimension can have just one Grand Total level. A Grand Total level does not contain dimensional attributes and does not have a level key. However, you can associate measures with a Grand Total level. The aggregation level for those measures will always be the grand total for the dimension.
- **Level.** All levels, except the Grand Total level, need to have at least one column. However, it is not necessary to explicitly associate all of the columns from a table with logical levels. Any column that you do not associate with a logical level is automatically associated with the lowest level in the dimension that corresponds to that dimension table. All logical columns in the same dimension table have to be associated with the same dimension.
- **Hierarchy.** Each dimension contains one or more hierarchies. All hierarchies must have a common leaf level and a common root (all) level.

For example, a time dimension might contain a fiscal hierarchy and a calendar hierarchy, with a common leaf level of Day. Day has two named parent levels called Fiscal Year and Calendar Year, which are both children of the All root level.

In the Business Model and Mapping layer, logical hierarchies are not defined as independent metadata objects, unlike hierarchies in the Presentation layer. Rather, logical hierarchies exist implicitly through the relationships between levels.

- **Level keys.** Each logical level (except the topmost level defined as a Grand Total level) must have one or more attributes that compose a level key. The level key defines the unique elements in each logical level. The dimension table logical key has to be associated with the lowest level of a dimension and has to be the level key for that level.

A logical level can have multiple level keys. When that is the case, specify the key that is the primary key of that level. All dimension sources which have an aggregate content at a specified level need to contain the column that is the primary key of that level. Each logical level should have one level key that is displayed when an Oracle BI Presentation Services user clicks to drill down. This may or may not be the primary key of the level. To set the level key to display, select the **Use for display** option in the Level Key dialog.

Be careful using level keys such as Month whose domain includes the values January, February, and so on (or in other words, values that are not unique to a

particular month, repeating every year). To define Month as a level key, you also need to include an attribute from a higher level (for example, Year). To add Year, click **Add** in this dialog and select the logical column from the dialog that is presented.

Level keys should be meaningful business keys (like Month_name='2010 July') rather than generated surrogate keys (like time_key='1023793'), because surrogate keys are physical artifacts that only apply to a single instance of a source table. The business name, in contrast, can map to any physical instance for that logical column. For example, month_name might map to a detailed table, an aggregate table from an aggregate star, and a column in a federated spreadsheet. Note that the Physical layer still uses the surrogate keys in the joins, so there is no performance or flexibility penalty for using business keys in the business model.

- **Time dimensions and chronological keys.** You can identify a dimension as a time dimension. At least one level of a time dimension must have a chronological key. The following is a list of some guidelines you should use when setting up and using time dimensions:
 - At least one level of a time dimension must have a chronological key. See ["Selecting and Sorting Chronological Keys in a Time Dimension"](#) for more information.
 - All time series measures using the AGO, TODATE, and PERIODROLLING functions must be on time levels. AGO, TODATE, and PERIODROLLING aggregates are created as derived logical columns. See ["About Time Series Functions"](#) for more information.
 - AGO, TODATE, and PERIODROLLING functionality is not supported either on fragmented dimensional logical table sources, or on fact sources fragmented on the same time dimension. Fact sources may be fragmented on other dimensions. See ["About Time Series Functions"](#) for more information.
- **Unbalanced (or ragged) hierarchy.** An unbalanced (or ragged) hierarchy is a hierarchy where the leaves (members with no children) do not necessarily have the same depth. For example, a site can choose to have data for the current month at the day level, previous months data at the month level, and the previous 5 years data at the quarter level.

User applications can use the ISLEAF function to determine whether to allow drilldown from any particular member. See ["ISLEAF"](#) for more information.

A missing member is implemented in the data source with a null value for the member value. All computations treat the null value as a unique child within its parent. Level-based measures and aggregate-by calculations group all missing nodes together.

Note that unbalanced hierarchies are not necessarily the same as parent-child hierarchies. Parent-child hierarchies are unbalanced by nature, but level-based hierarchies can be unbalanced also.


- **Skip-level hierarchy.** A skip-level hierarchy is a hierarchy where there are members that do not have a value for a particular ancestor level. For example, in a Country-State-City-District hierarchy, the city 'Washington, D.C.' does not belong to a State. In this case, you can drill down from the Country level (USA) to the City level (Washington, D.C.) and below.

In a query, skipped levels are not displayed, and do not affect computations. When sorted hierarchically, members appear under their nearest ancestors.

A missing member at a particular level is implemented in the data source with a null value for the member value. All computations treat the null value as a unique child within its parent. Level-based measures and aggregate-by calculations group all skip-level nodes together.

Figure 9–1 shows a hierarchy with both unbalanced and skip-level characteristics. For example, A-Brand 4, B-LOB 3, and Type 5 are unbalanced branches, while skips are present between A-Brand 2 and Type 3, B-LOB 2 and Product 6, and others.

Figure 9–1 Hierarchy with Unbalanced and Skip-Level Characteristics



Using Dimension Hierarchy Levels in Level-Based Hierarchies

Dimension hierarchical levels can be used to perform the following actions:

- Set up aggregate navigation
- Configure level-based measure calculations (refer to [Example 9–1](#)).
- Determine what attributes appear when Oracle BI Presentation Services users drill down in their data requests

Manually Creating Dimensions, Levels, and Keys with Level-Based Hierarchies

To create and manage dimension hierarchy levels in level-based hierarchies, perform the tasks described in the following sections:

- [Creating Dimensions in Level-Based Hierarchies](#)

- [Creating Logical Levels in a Dimension](#)
- [Associating a Logical Column and Its Table with a Dimension Level](#)
- [Identifying the Primary Key for a Dimension Level](#)
- [Selecting and Sorting Chronological Keys in a Time Dimension](#)
- [Adding a Dimension Level to the Preferred Drill Path](#)

Creating Dimensions in Level-Based Hierarchies

After creating a dimension, each dimension can be associated with attributes (columns) from one or more logical dimension tables and level-based measures from logical fact tables. After you associate logical columns with a dimension level, the tables in which these columns exist appear in the Tables tab of the Dimension dialog.

To create a dimension with a level-based hierarchy:

1. In the Business Model and Mapping layer of the Administration Tool, right-click a business model and select **New Object > Logical Dimension > Dimension with Level-Based Hierarchy**.

Note that this option is only available when there is at least one dimension table that has no dimension associated with it.

2. In the Logical Dimension dialog, in the General tab, type a name for the dimension.

The **Default root level** field is automatically populated after you associate logical columns with a dimension level.

3. If the dimension is a time dimension, select **Time**.
4. If the dimension is an unbalanced dimension, select **Ragged**.
5. If the dimension is a skip-level dimension, select **Skipped Levels**.

Note: It is a best practice to ensure that the physical hierarchy type set in the Physical layer matches the dimension properties you select in the Business Model and Mapping layer. See "[Working with Physical Hierarchy Objects](#)" for more information.

In addition, you must ensure that the Ragged and Skipped Levels dimension properties are set correctly for queries to work.

6. (Optional) Type a description of the dimension.
7. Click **OK**.

Creating Logical Levels in a Dimension

When creating logical levels in a dimension, you also create the hierarchy by identifying the type of level and defining child levels. See "[Automatically Creating Business Model Objects for Multidimensional Data Sources](#)" for more information about creating hierarchies for a multidimensional data source.

To define general properties for a logical level in a dimension:

1. In the Business Model and Mapping layer of the Administration Tool, right-click a dimension and select **New Object**, then select **Logical Level**.
2. In the Logical Level dialog, in the General tab, specify a name for the logical level.

3. For **Number of elements at this level**, specify the number of elements that exist at this logical level. If this level will be the Grand Total level, leave this field blank. The system will set to a value of 1 by default.

The number does not have to be exact, but ratios of numbers from one logical level to another should be accurate. For relational sources, you can retrieve the row count for the level key and use that number as the number of elements. For multidimensional sources, you can use the number of members at that level.

The Oracle BI Server uses this number when selecting which aggregate source to use. For example, when aggregate navigation is used, multiple fact sources exist at different grains. The Oracle BI Server multiplies the number of elements at each level for each qualified source as a way to estimate the total number of rows for that source. Then, the Oracle BI Server compares the result for each source and selects the source with the lowest number of total elements to answer the query. The source with the lowest number of total elements is assumed to be the fastest.

4. Choose one of the following options, if appropriate:
 - If the logical level is the Grand Total level, select **Grand total level**. There should be only one Grand Total level for a dimension.
 - If measure values at a particular level fully constitute aggregated measures at its parent level, select **Supports rollup to higher level of aggregation**.
5. To define child logical levels, click **Add**.
6. In the Browse dialog, select the child logical levels and click **OK**.

The child levels appear in the Child Levels pane.
7. To remove a previously defined child level, select the level in the Child Levels pane and click **Remove**.

The child level and all of its child levels are deleted from the Child Levels pane.
8. (Optional) Type a description of the logical level.
9. Click **OK**.


Associating a Logical Column and Its Table with a Dimension Level

After you create all logical levels within a dimension, you need to drag and drop one or more columns from the dimension table to each logical level except the Grand Total level. The first time you drag a column to a dimension it associates the logical table to the dimension. It also associates the logical column with that level of the dimension. To change the logical level to be associated with that logical column, you can drag a column from one logical level to another.

The logical column or columns that comprise the logical key of a dimension table must be associated with the lowest level of the dimension.

After you associate a logical column with a dimension level, the tables in which these columns exist appear in the Tables tab of the Dimensions dialog.

For time dimensions, ensure that all time-related logical columns in the source table are defined in the time dimension. For example, if a time-related logical table contains the columns Month Name and Month Code, you must ensure that both columns are dragged to the appropriate level within the dimension. [Figure 9-2](#) shows how to associate logical columns with a logical level.

Figure 9–2 Associating Logical Columns with a Logical Level**To verify tables that are associated with a dimension:**

1. In the Business Model and Mapping layer of the Administration Tool, double-click a dimension.
2. In the Dimensions dialog, click the Tables tab.
The tables list contains tables that you associated with that dimension. This list of tables includes only one logical dimension table and one or more logical fact tables (if you created level-based measures).
3. Click **OK** or **Cancel** to close the Dimensions dialog.

[Example 9–1](#) and [Example 9–2](#) show how to associate measures to different levels of level-based dimension hierarchies.

Example 9–1 Level-Based Measure Calculations

A level-based measure is a column whose values are always calculated to a specific level of aggregation. For example, a company might want to measure its revenue based on the country, based on the region, and based on the city. You can set up columns to measure CountryRevenue, RegionRevenue, and CityRevenue.

When a query includes a level-based measure column, and the query grain is higher than the level of aggregation specific to the column, the query results return null. Note that in previous releases, results were returned for this situation, but they were not deterministic.

The measure AllProductRevenue is an example of a level-based measure at the Grand Total level. Level-based measures allow a single query to return data at multiple levels of aggregation. They are also useful in creating share measures, that are calculated by taking some measure and dividing it by a level-based measure to calculate a percentage. For example, you can divide salesperson revenue by regional revenue to calculate the share of the regional revenue each salesperson generates.


To set up these calculations, you need to build a dimensional hierarchy in your repository that contains the levels Grandtotal, Country, Region, and City. This hierarchy contains the metadata that defines a one-to-many relationship between Country and Region and a one-to-many relationship between Region and City. For each country, there are many regions, but each region is in only one country. Similarly, for each region, there are many cities, but each city is in only one region.

Next, you need to create three logical columns (CountryRevenue, RegionRevenue, and CityRevenue). Each of these columns uses the logical column Revenue as its source. The Revenue column has a default aggregation rule of SUM and has sources in the underlying databases.

You then drag the CountryRevenue, RegionRevenue, and CityRevenue columns into the Country, Region, and City levels, respectively. Each query that requests one of these columns returns the revenue aggregated to its associated level.

Figure 9–3 shows what the business model in the Business Model and Mapping layer looks like for this example.

Figure 9–3 Example of Business Model in the Business Model and Mapping Layer



In the Geography Dimension, the CountryRevenue and RegionRevenue columns are attributes of the Country and Region levels. In the Sales Facts table, the Revenue column has a default aggregation rule of SUM and is mapped to physical detail data or physical aggregate data. CountryRevenue and RegionRevenue columns use the Revenue column as their source.

Example 9–2 Grand Total Dimension Hierarchy

You might have a product dimensional hierarchy with levels TotalProducts (Grand Total level), Brands, and Products. Additionally, there might be a column called Revenue that is defined with a default aggregation rule of Sum. You can then create a logical column, AllProductRevenue, that uses Revenue as its source (as specified in the General tab of the Logical Column dialog). Now, drag AllProductRevenue to the Grand Total level. Each query that includes this column returns the total revenue for all products. The value is returned regardless of any constraints on Brands or Products. If you have constraints on columns in other tables, the grand total is limited

to the scope of the query. For example, if the scope of the query asks for data from 1999 and 2000, the grand total product revenue is for all products sold in 1999 and 2000.

If you have three products, A, B, and C with total revenues of 100, 200, and 300 respectively, then the grand total product revenue is 600 (the sum of each product's revenue). If you have set up a repository as described in this example, the following query produces the results listed:

```
SELECT product, productrevenue, allproductrevenue
FROM sales_subject_area
WHERE product IN 'A' or 'B'
```

The results are as follows:

PRODUCT	PRODUCTREVENUE	ALLPRODUCTREVENUE
A	100	600
B	200	600

In this example, the AllProductRevenue column always returns a value of 600, regardless of the products on which the query constrains.

Identifying the Primary Key for a Dimension Level

Use the Keys tab in the Logical Level dialog to identify the primary key for a level.

To specify a primary key for a dimension level:

1. In the Business Model and Mapping layer of the Administration Tool, expand a dimension and then expand the highest level (Grand Total level) of the dimension.
2. Double-click a logical level below the Grand Total level.
3. In the Logical Level dialog, click the Keys tab.
4. In the Keys tab, from the **Primary key** list, select a level key.
If only one level key exists, it is the primary key by default.
5. To add a column to the list, perform the following steps:
 - a. In the Logical Level dialog, click **New**.
 - b. In the Logical Level Key dialog, type a name for the key.
 - c. In the Logical Level Key dialog, select a column or click **Add**.
 - d. If you click **Add**, in the Browse dialog, select the column, and then click **OK**.

The column you selected appears in the **Columns** list of the Logical Level Key dialog and is automatically selected.

Note: You cannot use a derived logical column that is the result of a LOOKUP function as part of a primary logical level key. This limitation exists because the LOOKUP operation is applied after aggregates are computed, but level key columns must be available before the aggregates are computed because they define the granularity at which the aggregates are calculated.

You can use a derived logical column that is the result of a LOOKUP function as a secondary logical level key.

6. If the level is in a time dimension, you can select chronological keys and sort the keys by name.

7. (Optional) Type a description for the key and then click **OK**.
8. Repeat Step 2 through Step 7 to add primary keys to other logical levels.
9. In the Logical Level dialog, click **OK**.

Selecting and Sorting Chronological Keys in a Time Dimension

At least one level of a time dimension must have a chronological key. Although you can select one or more chronological keys for any level and then sort keys in the level, only the first chronological key is used.

To select and sort chronological keys for a time dimension:

1. In the Business Model and Mapping layer of the Administration Tool, expand a time dimension and then expand the highest level (Grand Total level) of the dimension.

Note: For a dimension to be recognized as a time dimension, you must select **Time** on the General tab of the Dimension dialog.

2. Double-click a logical level below the Grand Total level.
3. In the Logical Level dialog, click the Keys tab.
4. To select a chronological key, in the Keys tab, select the **Chronological Key** option. You may need to scroll to the right to see this option.
5. To sort chronological keys, in the Keys tab, double-click a chronological key.
6. In the Chronological Key dialog, select a chronological key column, click **Up** or **Down** to reorder the column, and then click **OK**.

Adding a Dimension Level to the Preferred Drill Path

You can use the Preferred Drill Path tab to identify the drill path to use when Oracle BI Presentation Services users drill down in their data requests. You should use this only to specify a drill path that is outside the normal drill path defined by the dimensional level hierarchy. It is most commonly used to drill from one dimension to another. You can delete a logical level from a drill path or reorder a logical level in the drill path.

To add a dimension level to the preferred drill path:

1. Click **Add** to open the Browse dialog, then select the logical levels to include in the drill path. You can select logical levels from the current dimension, or from other dimensions.
2. Click **OK** to return to the Level dialog.

The names of the levels are added to the Names pane.

Automatically Creating Dimensions with Level-Based Hierarchies

You can set up a dimension automatically from a logical dimension table if a dimension for that table does not exist. To create a dimension automatically, the Administration Tool examines the logical table sources and the column mappings in those sources and uses the joins between physical tables in the logical table sources to determine logical levels and level keys. Therefore, it is best to create a dimension in this way after all the logical table sources have been defined for a dimension table.

The following rules apply:

- Create Dimensions is only available if the selected logical table is a dimension table (defined by 1:N logical joins) and no dimension has been associated with this table.
- An automatically created dimension uses the same name as the logical table, adding Dim as a suffix. For example, if a table is named Periods, the dimension is named Periods Dim.
- A Grand Total level is automatically named *logical_table_name* Total. For example, the Grand Total level of the Periods Dim table is Periods Total.
- When there are multiple tables in a source, the join relationships between tables in the source determine the physical table containing the lowest-level attributes. The lowest level in the hierarchy is named *logical_table_name* Detail. For example, the lowest level of the periods table is Periods Detail.
- The logical key of the dimension table is mapped to the lowest level of the hierarchy and specified as the level key. This logical column should map to the key column of the lowest level table in the dimension source.
 - If there are two or more physical tables in a source, the columns that map to the keys of those tables become additional logical levels. These additional level names use the logical column names of these key columns.
 - The order of joins determines the hierarchical arrangement of the logical levels. The level keys of these new logical levels are set to the logical columns that map to the keys of the tables in the source.
- If there are multiple logical table sources, the tool uses attribute mappings and physical joins to determine the hierarchical order of the tables in the physical sources. For example, you might have three sources (A, B, C) each containing a single physical table and attribute mappings for 10, 15, and 3 attributes, respectively (not counting columns that are constructed from other logical columns). The following is a list of the results of creating a dimension for this table automatically:
 - The Administration Tool creates a dimension containing four logical levels, counting the Grand Total and detail levels.
 - The key of the table in source B (that has the greatest number of columns mapped and contains the column mapping for the logical table key) should be the level key for the detail level.
 - The parent of the detail level should be the logical level named for the logical column that maps to the key of the physical table in source A.
 - Any attributes that are mapped to both A and B should be associated with level A.
 - The parent of level A should be the logical level named for the logical column that maps to the key of the physical table in source C.
 - Any columns that are mapped to both A and C should be associated with level C.
- Table joins in a physical source might represent a pattern that results in a split hierarchy. For example, the Product table can join to the Flavor table and a Subtype table. This would result in two parents of the product detail level, one flavor level and one subtype level.
- You cannot create a dimension automatically in the following situations:

- If a dimension with joins and levels has already been created, Create Dimension does not appear on the right-click menu.
- If the table is not yet joined to any other table, the option is not available because it is considered a fact table.
- In a snowflake schema, if you use a table with only one source and create the dimension automatically, the child tables are automatically incorporated into a hierarchy. The child tables form intermediate levels between the Grand Total level and detail level. If more than one child table exists for a dimension table, the hierarchy is a split hierarchy.

To create a dimension automatically:

1. In the Administration Tool, open a repository.
2. In the Business Model and Mapping layer, right-click a logical dimension table that is not associated with any dimension .
3. From the right-click menu, select **Create Logical Dimension**, then select either **Dimension with Level-Based Hierarchy** or **Dimension with Parent-Child Hierarchy**.

The new dimension is displayed in the Business Model and Mapping layer.

Populating Logical Level Counts Automatically

Estimate Levels enables administrators to automatically populate level counts for one or more dimension hierarchies. Level counts are utilized by the query engine to determine the most optimal query plan and optimizes overall system performance.

The repository must be opened in online mode and the business model must be available for queries. Then, in the Business Model and Mapping layer, you can select any of the following logical layer elements, and then execute the Estimate Levels command:

- **Business model.** Must be available for queries. If you select this object, the Administration Tool attempts to check out all objects in the business model.
- **Dimension.** You should run a consistency check on dimensions to ensure that the dimension is logically sound.
- **A combination of business models and dimensions.** You can select multiple dimensions and multiple business models individually.

When run, the Estimate Levels command also launches a consistency check on the level counts as described in the following list:

- Checks that a level key is valid. Columns in levels have referential integrity.
- Checks the parent-child relationship. If the parent level count is greater than the child level count, an error is returned.
- Generates a run report that lists all the counts that were estimated and any errors or consistency warnings.
- The queries and errors are logged to nquery.log on the Oracle BI Server.

Set the log level at 4 or higher to write this information to the log file. For more information about logging, see "Diagnosing and Resolving Issues in Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

To populate logical level counts automatically:

1. In the Administration Tool, open a repository in online mode.
2. Right-click one or more business models and dimension objects, and select **Estimate Levels**.
3. In the Check Out Objects dialog, click **Yes** to check out the objects that appear in the list.

If you click **No**, the action terminates because you must check out items to run Estimate Levels.

In the Administration Tool dialog, a list of the dimension level counts and any errors or warning messages appear.

When you check in the objects, you can check the global consistency of your repository.

Creating and Managing Dimensions with Parent-Child Hierarchies

A parent-child hierarchy is a hierarchy of members that all have the same type. This contrasts with level-based hierarchies, where members of the same type occur only at a single level of the hierarchy.

This section contains the following topics:

- [About Parent-Child Hierarchies](#)
- [Creating Dimensions with Parent-Child Hierarchies](#)
- [Defining Parent-Child Relationship Tables](#)

About Parent-Child Hierarchies

The most common real-life occurrence of a parent-child hierarchy is an organizational reporting hierarchy chart, where the following all apply:


- Each individual in the organization is an employee.
- Each employee, apart from the top-level managers, reports to a single manager.
- The reporting hierarchy has many levels.

These conditions illustrate the basic features that define a parent-child hierarchy, namely:

- A parent-child hierarchy is based on a single logical table (for example, the "Employees" table)
- Each row in the table contains two identifying keys, one to identify the member itself, the other to identify the "parent" of the member (for example, Emp_ID and Mgr_ID)

[Figure 9-4](#) shows an example of a multi-level parent-child hierarchy.

Figure 9–4 Multi-Level Parent-Child Hierarchy



The following table shows how this parent-child hierarchy could be represented by the rows and key values in an Employees table.

Emp_ID	Mgr_ID
Andrew	<i>null</i>
Barbara	Andrew
Carlos	Andrew
Dawn	Barbara
Emre	Barbara

You can expose logical dimensions with parent-child hierarchies to Oracle BI Answers users by creating presentation hierarchies that are based on particular logical dimensions. Creating hierarchies in the Presentation layer enables users to create hierarchy-based queries. See ["Working with Presentation Hierarchies and Levels"](#) for more information.

This section contains the following topics:

- [About Levels and Distances in Parent-Child Hierarchies](#)
- [About Parent-Child Relationship Tables](#)
- [About Parent-Child Hierarchies Populated with Preaggregated Data](#)

About Levels and Distances in Parent-Child Hierarchies

Unlike the situation with level-based hierarchies, all the dimension members of a parent-child hierarchy occur in a single logical column. In a parent-child hierarchy, the parent of a member is in another row in the same logical column, pointed to by the parent key. This is unlike a level-based hierarchy, where the parent of a member is in a different logical column in the same row. In other words, navigation in a parent-child hierarchy follows data values, while navigation in a level-based hierarchy follows the metadata structure.

In level-based hierarchies, each level is named, and occupies a position in the hierarchy that corresponds to a real-world attribute or category that is deemed useful for analysis. Unlike level-based hierarchies, where the number of levels is fixed at design time, there is no limit to the depth of a parent-child hierarchy, and the depth can change at run time due to new data."

Every Oracle BI Server parent-child hierarchy has two system-generated entities, "Total" and "Detail," that are automatically defined for each parent-child hierarchy

when the logical dimension is created. The "Detail" entity contains all the hierarchy members.

These two system-generated entities are different from the implicit, inter-member levels between ancestors and descendants in a parent-child hierarchy. These implicit levels are referred to as parent-child hierarchical levels in this section.

Closely associated with levels is the concept of distance in parent-child hierarchies. The distance of one member from another is the number of parent-child hierarchical levels from the member to an ancestor or to a descendant. For example, the distance from a member to its parent is always 1, and the distance from Andrew to Emre in [Figure 9-4](#) is 2.

About Parent-Child Relationship Tables

In relational tables, the relationships between different members in a parent-child hierarchy are implicitly defined by the identifier key values in the associated base table.

However, for each Oracle BI Server parent-child hierarchy defined on a relational table, you must also explicitly define the inter-member relationships in a separate parent-child relationship table.

The parent-child relationship table must include four columns, as follows:

- A column that identifies the member
- A column that identifies an ancestor of the member

Note: The ancestor may be the parent of the member, or a higher-level ancestor.

- A "distance" column that specifies the number of parent-child hierarchical levels from the member to the ancestor
- A "leaf" column that indicates if the member is a leaf member (1=Yes, 0=No)

The column names can be user defined. The data types of the columns must satisfy the following conditions:

- The member and ancestor identifier columns have the same data type as the associated columns in the logical table that contains the hierarchy members.

Note that the example shown in [Table 9-1](#) uses text strings for readability, but you normally use integer surrogate keys for member_key and ancestor_key, if they exist in the source dimension table.

- The "distance" and "leaf" columns are INTEGER columns.

Note the following about the rows in a parent-child relationship table:

- Each member must have a row pointing at itself, with distance zero.
- Each member must have a row pointing at each of its ancestors. For a root member, this is a termination row with null for the parent and distance values.

[Table 9-1](#) shows an example of a parent-child relationship table with rows that represent the inter-member relationships for the employees shown in [Figure 9-4](#).

Table 9–1 Example Parent-Child Relationship Table

Member_Key	Ancestor_Key	Distance	Isleaf
Andrew	Andrew	0	0
Barbara	Barbara	0	0
Carlos	Carlos	0	0
Dawn	Dawn	0	0
Emre	Emre	0	0
Andrew	null	null	0
Barbara	Andrew	1	0
Carlos	Andrew	1	1
Dawn	Barbara	1	1
Dawn	Andrew	2	1
Emre	Barbara	1	1
Emre	Andrew	2	1

Typically, you generate scripts to create and populate the parent-child relationship table through a wizard that you can invoke when you define the parent-child hierarchy. Note the following about the create and load scripts:

- You run the create script only once, to create the parent-child relationship table in the data source.
- You must run the load script after each time the data changes in the dimension table. Because of this, you typically call the load script in your ETL processing. The load script reloads the entire parent-child relationship table; it is not incremental.

If you do not choose to use the wizard, then you must create the parent-child relationship table manually and then import it into the Physical layer before associating it with the parent-child hierarchy. In this latter case, it is also your responsibility to populate the table with the data required to describe the inter-member relationships in the parent-child hierarchy.

About Parent-Child Hierarchies Populated with Preaggregated Data

Some parent-child hierarchies contain preaggregated data that is populated for all nodes of the hierarchy. For example, a root node might be populated with the aggregation of the data for all of its descendent nodes. It is important to ensure that queries do not aggregate the members from this dimension to avoid erroneous results.

To correctly model this type of parent-child hierarchy, you must still create a parent-child relationship table to support hierarchical filter functions like IsAncestor and IsDescendant. However, you can join the parent-child dimension table directly with the fact table rather than joining through the parent-child relationship table. Doing so ensures that the preaggregated member value is returned, rather than rolling up all the descendants.

Note: Do not modify the parent-child relationship table script to only include the "self" rows, because doing so would break the IsAncestor and IsDescendant functions.

To achieve the correct aggregation for dimensions of this type, you must first determine what you want to see as a grand total when the parent-child hierarchy is aggregated. For example, assume that your hierarchy contains a single root member, and you want to display the preaggregated value for this root member. In this case, you first create an additional fact logical table source mapped at the Total level of the parent-child hierarchy. Then, in the logical table source, create a WHERE clause filter that selects only the root member.

With this model in place, for queries that are at the Total level of the parent-child hierarchy, the Oracle BI Server selects the aggregate logical table source and applies the root member WHERE clause filter. For queries that are at the Detail level, the Oracle BI Server selects the detailed logical table source and returns the preaggregated member values. In either case, it does not matter how the aggregation rule is set, because there is a preaggregated source at each level.

Note that this approach only works if the queries are either at the Total or Detail level of the parent-child dimension. However, for queries that group by some non-unique attribute of the parent-child dimension, the aggregation might not be correct. For example, if an Employee dimension has a Location attribute, and a query groups by Employee.Location, then there will likely be some double counting because an employee often reports to other employees at the same location. Because of this, when fact tables contain preaggregated member values, you should avoid grouping by non-unique attributes of the parent-child dimension. If grouping by those attributes is unavoidable, then you should model them as separate dimensions.

Creating Dimensions with Parent-Child Hierarchies

The key elements that you must define for a parent-child hierarchy are the identifier columns for the member and the parent of the member. This basic principle applies to all parent-child hierarchies, regardless of the data source from which the hierarchy is derived.

Parent-child hierarchies that are based on relational tables must have an accompanying parent-child relationship table. See ["About Parent-Child Relationship Tables"](#) for more information.

To create dimensions with a parent-child hierarchy:

1. In the Business Model and Mapping layer of the Administration Tool, perform one of the following steps:
 - Right-click a business model and select **New Object > Logical Dimension > Dimension with Parent-Child Hierarchy**. Note that this option is only available if there is at least one logical dimension table in the business model that has no dimension associated with it.
 - Right-click a dimension table that is not associated with any dimension and select **Create Logical Dimension**, then select **Dimension with Parent-Child Hierarchy**.
2. In the Logical Dimension dialog, in the General tab, type a name for the dimension.
3. Click **Browse** beside the Member Key field.

The Browse window shows the logical dimension tables in the business model, each with their primary keys.
4. Select a Member Key for the parent-child hierarchy and click **OK**.
5. Click **Browse** beside the Parent Key field.

The Browse window shows the columns, other than the primary key, in the logical table that you selected in step 4.

6. Select a column that will be the Parent Key for the parent-child hierarchy and click **OK**.
7. If the logical table that you selected in step 4 is not from a relational table source, click **OK** to finish the process of creating the dimension.

If the logical table you selected in step 4 is from a relational table source, you must continue the dimension definition process by setting up the parent-child relationship table for the hierarchy. See ["Defining Parent-Child Relationship Tables"](#) for details.

Defining Parent-Child Relationship Tables

For parent-child hierarchies based on relational tables, you must define a parent-child relationship table. See ["About Parent-Child Relationship Tables"](#) for more information.

When you create the parent-child relationship table, you must choose one of the following options:

- Select a previously-created parent-child relationship table
- Use a wizard that will generate scripts to create and populate the parent-child relationship table

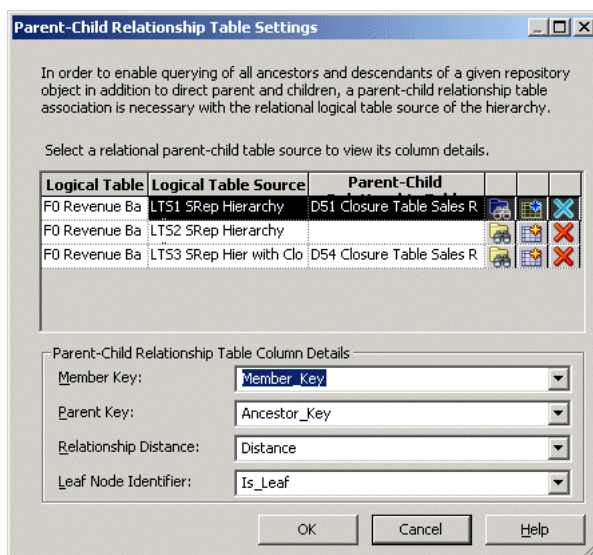
To define parent-child relationship tables:

1. In the Logical Dimension dialog, click **Parent-Child Settings**.

The Parent-Child Relationship Table Settings windows appears, with the Logical Table and Logical Table Source values filled in.

[Figure 9–5](#) shows the Parent-Child Relationship Table Settings dialog.

Figure 9–5 Parent-Child Relationship Table Settings Dialog



2. You can either manually define the parent-child relationship table for the hierarchy, or you can start a wizard that will perform the definition for you (recommended).
 - To start the manual process, continue at step 3.

- To start the wizard, continue at step 7.
3. Click the **Select Parent-Child Relationship Table** button to start the manual process of defining the parent-child relationship table for the parent-child hierarchy.
4. Select the physical table that acts as the parent-child relationship table for your parent-child hierarchy. The table must already exist in the Physical layer.

The parent-child relationship table must have at least four columns that describe how the inter-member relationships are derived in the logical table selected for the hierarchy. See "[About Parent-Child Relationship Tables](#)" for more information.
5. Map the four columns from the physical parent-child relationship table to the fields in the Parent-Child Table Column Details area, as follows:
 - Select the Member Key column
 - Select the Parent Key column
 - Select the Relationship Distance column
 - Select the Leaf Node Identifier column
6. Click **OK**, then click **OK** again to finish the manual process of defining the parent-child relationship table.
7. Click the **Create Parent-Child Relationship Table** button to start the wizard.

The Generate Parent-Child Relationship Table Wizard generates SQL scripts for creating and populating the parent-child relationship table. At the end of the wizard, the Oracle BI Server stores the scripts into directories chosen during the wizard session. The scripts, when executed, will make the parent-child relationship table available to the metadata repository.

The wizard contains the following three main windows:

 - Script Location
 - Parent-Child Relationship Table Details
 - Preview Script
8. In the Generate Parent-Child Relationship Table - Script Location screen, enter the **Name** for the **DDL Script to Generate the Parent-Child Table**, and select the **Location** where the Generate script will be placed.
9. Enter the **Name** for the **DDL Script to Populate the Parent-Child Table**, and select the **Location** where the Populate script will be placed.
10. Click **Next**.
11. In the Parent-Child Relationship Table Details screen, enter the **Name** for the parent-child relationship table.
12. Click **Browse** beside the Catalog/Schema field to select the catalog or schema for the parent-child relationship table.
13. Click **Next**.
14. In the Preview Script window, you can view either or both of the scripts.
15. Click **Finish**.
16. In the Parent-Child Relationship Table Settings window, click **OK**.
17. In the Logical Dimension window, click **OK**.

18. If you used the Generate Parent-Child Relationship Table Wizard to generate create and load scripts, run the scripts to create and load the parent-child relationship table in your data source.

Adding the Parent-Child Relationship Table to the Model

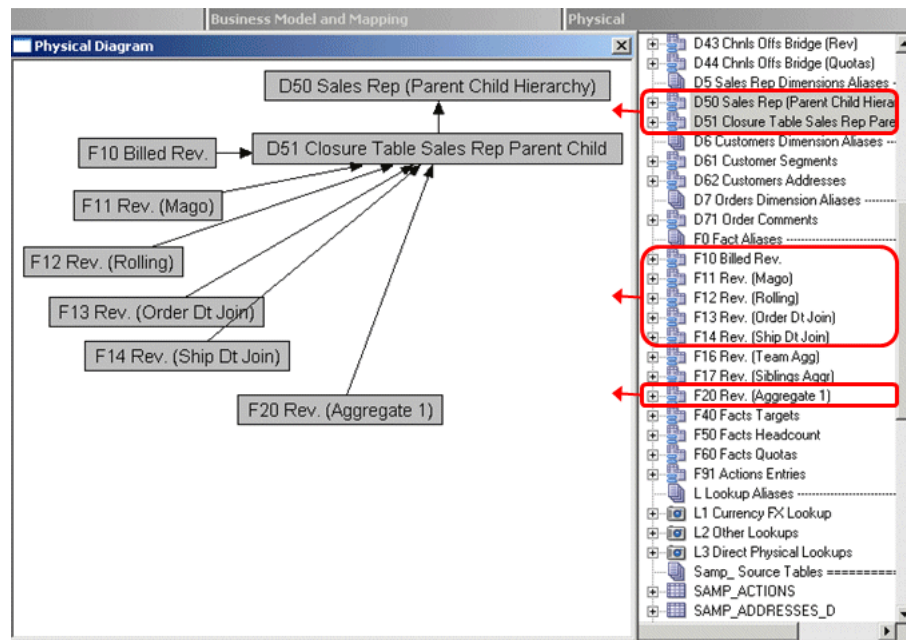
After you have created a parent-child relationship table and imported it into the Physical layer (either manually or using the Generate Parent-Child Relationship Table Wizard), you must edit Physical layer joins to include the parent-child relationship table. You also need to add the parent-child relationship table to the appropriate logical table source.

To add the parent-child relationship table to the model:

1. In the Administration Tool, in the Physical layer of the repository, open the Physical Diagram so that it shows the parent-child relationship table and associated dimension table and fact tables. To do this, right-click the appropriate physical tables and select **Physical Diagram > Selected Object(s) Only**.
2. Delete the direct joins from the dimension table to each of the fact tables.
3. Create joins from each the fact tables to the dimension table through the parent-child closure table, as follows:
 - a. Create a join from the parent-child relationship table to the dimension table using the ancestor key.
 - b. Create joins from the fact tables to the parent-child relationship table using the member key.

Figure 9–6 shows joins from a dimension table to fact tables that go through a parent-child relationship table.

Figure 9–6 Physical Layer Joins Through a Parent-Child Relationship Table



4. In the Business Model and Mapping layer, double-click the logical table source for the logical dimension table that is used in your parent-child hierarchy.
5. In the General tab of the Logical Table Source dialog, click the **Add** button.

6. Browse to locate the parent-child relationship table in the Physical layer and click **Select**.
7. Click **OK** in the Logical Table Source dialog.

Maintaining Parent-Child Hierarchies Based on Relational Tables

For parent-child hierarchies based on relational tables, you must ensure that the data in the parent-child relationship table accurately reflects the inter-member relationships in the dimension.

You may have created scripts manually to create and populate the parent-child relationship table, or you may have used the Generate Parent-Child Relationship Table Wizard to create the scripts. You must run these scripts, adapting them if necessary, as often as required to guarantee the integrity of the parent-child relationships in the hierarchy. You typically want to add the Populate script to your ETL process so that it runs after the dimension table is updated.

Modeling Time Series Data

Time series functions provide the ability to compare business performance with previous time periods, allowing you to analyze data that spans multiple time periods. For example, time series functions enable comparisons between current sales and sales a year ago, a month ago, and so on.

Because SQL does not provide a direct way to make time comparisons, you must model time series data in the Oracle BI repository. First, set up time dimensions based on the period table in your data warehouse. Then, you can define measures that take advantage of this time dimension to use the `AGO`, `TODATE`, and `PERIODROLLING` functions. At query time, the Oracle BI Server then generates highly optimized SQL that pushes the time offset processing down to the database whenever possible, resulting in the best performance and functionality.

This section contains the following topics:

- [About Time Series Functions](#)
- [Creating Logical Time Dimensions](#)
- [Creating `AGO`, `TODATE`, and `PERIODROLLING` Measures](#)

About Time Series Functions

Time series functions operate on time-oriented dimensions. To use these functions on a particular dimension, you must designate the dimension as a Time dimension and set one or more keys at one or more levels as chronological keys. These keys identify the chronological order of the members within a dimension level.

Time series functions include `AGO`, `TODATE`, and `PERIODROLLING`. These functions let you use Expression Builder to call a logical function to perform time series calculations instead of aliasing physical tables and modeling logically. The time series functions calculate `AGO`, `TODATE`, and `PERIODROLLING` functions based on the calendar tables in your data warehouse, not on standard SQL date manipulation functions.

[Figure 9-7](#) shows a sample report that includes several measures derived using time series functions.

Figure 9–7 Example Measures Derived Using Time Series Functions

	2008 Q1			2008 Q2			2008 Q3			2008 Q4		
	2008 / 01	2008 / 02	2008 / 03	2008 / 04	2008 / 05	2008 / 06	2008 / 07	2008 / 08	2008 / 09	2008 / 10	2008 / 11	2008 / 12
Dollars	100	200	300	101	202	303	110	220	330	444	555	666
Dollars Qago				100	200	300	101	202	303	110	220	330
Dollars QTD	100	300	600	101	303	606	110	330	660	444	999	1,665
Dollars 3-Period Rolling Sum	100	300	600	601	603	606	615	633	660	994	1,329	1,665
Dollars 3-Period Rolling Avg	33.3	100.0	200.0	200.3	201.0	202.0	205.0	211.0	220.0	331.3	443.0	555.0

Several different grains may be used in the time query, such as:

- **Query grain.** The lowest time grain of the request. In the report example shown in [Figure 9–7](#), the query grain is Month.
- **Time Series grain.** The grain at which the aggregation or offset is requested, for both AGO and TODATE functions. In the report example shown in [Figure 9–7](#), the time series grain is Quarter. Time series queries are valid only if the time series grain is at the query grain or longer. Note that the PERIODROLLING function does not have a time series grain; instead, you specify a start and end period in the function.
- **Storage grain.** The report example shown in [Figure 9–7](#) can be computed from daily sales or monthly sales. The grain of the source is called the storage grain. A chronological key must be defined at this level for the query to work, but performance is generally much better if a chronological key is also defined at the query grain.

Note that queries against time series data must be an exact match to hit the query cache. See *Oracle Fusion Middleware System Administrator’s Guide for Oracle Business Intelligence Enterprise Edition* for more information about the Oracle BI Server query cache.

The following sections describe the time series conversion functions:

- [About the AGO Function](#)
- [About the TODATE Function](#)
- [About the PERIODROLLING Function](#)

About the AGO Function

The AGO function offsets the time dimension to display data from a past period. This function is useful for comparisons, such as *Dollars* compared to *Dollars a Quarter Ago*. Note that the value of "Dollars Qago" for month 2008/08 equals the value of "Dollars" for month 2008/05.

[Figure 9–8](#) shows example values for the Dollars and Dollars Qago measures.

Figure 9–8 Example Dollars and Dollars Qago Measures

	2008 Q1			2008 Q2			2008 Q3			2008 Q4		
	2008 / 01	2008 / 02	2008 / 03	2008 / 04	2008 / 05	2008 / 06	2008 / 07	2008 / 08	2008 / 09	2008 / 10	2008 / 11	2008 / 12
Dollars	100	200	300	101	202	303	110	220	330	444	555	666
Dollars Qago				100	200	300	101	202	303	110	220	330

In the example shown in [Figure 9–8](#), the Dollars Qago measure is derived from the Dollars measure.

In Expression Builder, the AGO function has the following template:


```
Ago(<<Measure>>, <<Level>>, <<Number of Periods>>)
```

<<Measure>> represents the logical measure column from which you want to derive. In this example, you would select the measure "Dollars" from your existing logical fact tables.

<<Level>> is the optional time series grain you want to use. In this example, you would select "Quarter" from your time dimension.

<<Number of Periods>> is the size of the offset, measured in the grain you provided in the <<Level>> argument. For example, if the <<Level>> is Quarter and the <<Number of Periods>> is 2, the function displays dollars from two quarters ago.

Using this function template, you can create an expression for a One Quarter Ago measure, as follows:

```
Ago("Sales"."Base Measures"."Dollars" , "Sales"."Time MonthDim"."Quarter" , 1)
```

The <<Level>> parameter is optional. If you do not want to specify a time series grain in the AGO function, the function uses the query grain as the time series grain.

For example, you could define Dollars_Ago as Ago(Dollars, 1). Then, you could perform the following logical query:

```
SELECT Month, Dollars, Dollars_Ago
```

The result is the same as if you defined Dollars_Ago as Ago(Dollars, Month, 1). Alternatively, you could perform the following logical query:

```
SELECT Quarter, Dollars, Dollars_Ago
```

The result is the same as if you defined Dollars_Ago as Ago(Dollars, Quarter, 1).

See ["AGO"](#) for additional information about the AGO function syntax.

About the TODATE Function

The TODATE function accumulates the measure from the beginning of the time series grain period to the current displayed query grain period. For example, [Figure 9-9](#) shows a report with the measure "Dollars QTD," which is the Quarter To Date version of the "Dollars" measure.

Figure 9-9 Example Dollars and Dollars QTD Measures

	2008 Q1			2008 Q2			2008 Q3			2008 Q4		
	2008 / 01	2008 / 02	2008 / 03	2008 / 04	2008 / 05	2008 / 06	2008 / 07	2008 / 08	2008 / 09	2008 / 10	2008 / 11	2008 / 12
Dollars	100	200	300	101	202	303	110	220	330	444	555	666
Dollars QTD	100	300	600	101	303	606	110	330	660	444	999	1,665

In the example shown in [Figure 9-9](#), Dollars QTD for Month 2008/05 is the sum of Dollars for 2008/04 and 2008/05. In other words, Dollars QTD is the sum of the values for all the query grain periods (month) for the current time series grain period (quarter). The accumulation starts over for the next quarter.

In the example shown in [Figure 9-9](#), the Dollars QTD measure is derived from the Dollars measure.

In Expression Builder, the TODATE function has the following template:

```
ToDate(<<Measure>>, <<Level>>)
```

<<Measure>> represents the logical measure column from which you want to derive. In this example, you would select the measure "Dollars" from your existing logical fact tables.

<<Level>> is the time series grain you want to use. In this example, you would select "Quarter" from your time dimension.

Using this function template, you can create the following expression for the measure:

```
ToDate("Sales"."Base Measures"."Dollars" , "Sales"."Time MonthDim"."Quarter" )
```

Note that the query grain is specified in the query itself at run time. For example, this measure can display Quarter To Date at the Day grain, but still accumulates up to the end of the Quarter.

See ["TODATE"](#) for additional information about TODATE function syntax.

About the PERIODROLLING Function

The PERIODROLLING function lets you perform an aggregation across a specified set of query grain periods, rather than within a fixed time series grain. The most common use is to create rolling averages, such as "13-week Rolling Average."

Note that because this function has no time series grain, the length of the rolling sequence is determined by the query grain. For example, "Dollars 3-Period Rolling Average" averages the last 3 months if the query grain is Month, but averages the last 3 years if the query grain is Year.

This section describes how to build two measures using the PERIODROLLING function: "Dollars 3-Period Rolling Sum," and "Dollars 3-Period Rolling Average." [Figure 9–10](#) shows a report with these two measures.

Figure 9–10 Example Dollars, Dollars 3-Period Rolling Sum, and Dollars 3-Period Rolling Avg Measures

	2008 Q1			2008 Q2			2008 Q3			2008 Q4		
	2008 / 01	2008 / 02	2008 / 03	2008 / 04	2008 / 05	2008 / 06	2008 / 07	2008 / 08	2008 / 09	2008 / 10	2008 / 11	2008 / 12
Dollars	100	200	300	101	202	303	110	220	330	444	555	666
Dollars 3-Period Rolling Sum	100	300	600	601	603	606	615	633	660	994	1,329	1,665
Dollars 3-Period Rolling Avg	33.3	100.0	200.0	200.3	201.0	202.0	205.0	211.0	220.0	331.3	443.0	555.0

In the example shown in [Figure 9–10](#), the Dollars 3-Period Rolling Sum and Dollars 3-Period Rolling Avg measures are derived from the Dollars measure.

In Expression Builder, the PERIODROLLING function has the following template:

```
PeriodRolling(<<Measure>>, <<Starting Period Offset>>, <<Ending Period Offset>>)
```

<<Measure>> represents the logical measure column from which you want to derive. To create the measure Dollars 3-Period Rolling Sum, you would select the measure "Dollars" from your existing logical fact tables.

<<Starting Period Offset>> and <<Ending Period Offset>> identify the first period and last period used in the rolling aggregation, respectively. The integer is the relative number of periods from the displayed period. In the example shown in [Figure 9–10](#), the query grain is month, and the 3-month rolling sum starts 2 periods in the past and includes the current period. That is, for month 2008/07, the rolling sum includes 2008/05, 2008/06 and 2008/07. Therefore, to create the measure Dollars 3-Period Rolling Sum, the integers to indicate these offsets are -2 and 0.

Using this function template, you can create the following expression for the measure:

```
PeriodRolling("Sales"."Base Measures"."Dollars" , -2, 0)
```

The example shown in [Figure 9–10](#) also shows a 3-month rolling average. To compute this measure, you can divide the rolling sum that you previously created by 3 to get a 3-period rolling average. We know to divide the rolling sum by 3 because the <<Starting Period Offset>> and <<Ending Period Offset>> fields for the rolling sum are -2 and 0.

The expression for the 3-month rolling average is:

```
PeriodRolling("Sales"."Base Measures"."Dollars" , -2, 0) /3
```

It is usually a mistake to use the AVG function to create a rolling average. AVG computes the average of the database rows accessed at the storage grain, but you need an average where the denominator is the number of rolling periods at the query grain.

Note that the PERIODROLLING function includes a fourth optional hierarchy argument that lets you specify the name of a hierarchy in a time dimension, such as *yr*, *mon*, *day*, that you want to use to compute the time window. This option is useful when there are multiple hierarchies in a time dimension, or when you want to distinguish between multiple time dimensions. See "[PERIODROLLING](#)" for more information about the hierarchy argument and for details on the function syntax.

Creating Logical Time Dimensions

Compared to modeling an ordinary dimension, the time dimension requires just two additional steps: selecting the **Time** option in the Logical Dimension dialog, and designating a chronological key for every level of every dimension hierarchy.

Follow these additional guidelines when modeling time series data:

- It only makes sense to use a time series function when the data source contains history. Usually, a relational database that contains history uses a star or snowflake schema with an explicit time dimension table. A normalized, historical database is much rarer, but would still include a time hierarchy with levels in a schema similar to a snowflake. A simple date field is not adequate.
- Oracle Business Intelligence requires the time dimension physical table (or set of normalized tables) to be separate from the physical fact table to which it is related.

However, a somewhat common source schema pattern is a fully denormalized relational table or flat file, where the time dimension columns are in the same table as the facts and other dimensions. This cannot qualify as a time dimension, because the time dimension table is combined with the fact table.

In this case, if you cannot change the model in the source, the best practice is to create an Opaque View of the physical table containing the time columns, which acts as the distinct physical time dimension table. This Opaque View time dimension must then be joined to the physical table that contains the facts.

- In the Physical layer, the time dimension table (or lowest-level table in the normalized/snowflake case) must join directly to the fact table without intervening tables. This join must be a foreign key join.
- The tables in the physical model containing the time dimension cannot join to other data sources, except at the most detailed level.
- A member value (for example, a row in relational sources) must be physically present for every period at every hierarchy level. There cannot be any skips in the

sequence. Note that it does not matter whether there is fact data for every period; only the dimension data must be complete.

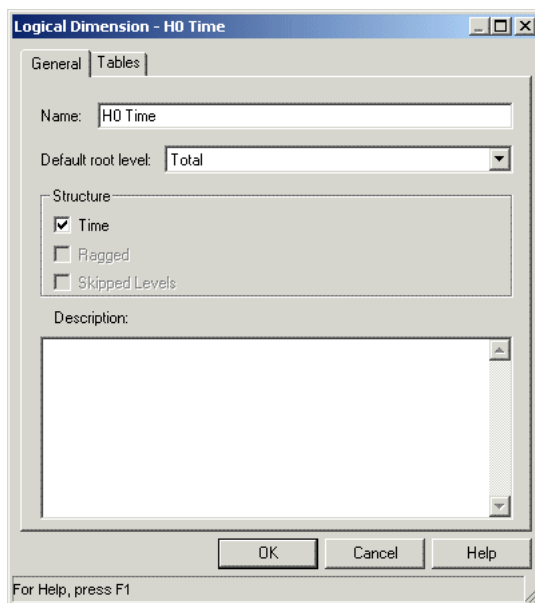
- Each unit of distance between members, such as "month," "half," or "year," must be modeled in a separate hierarchy level.

Selecting the Time Option in the Logical Dimension Dialog

Select the **Time** option in the General tab of the Logical Dimension dialog to enable time series functions on this dimension. Only logical dimensions with the **Time** option selected can be used as the time dimension for the time series functions AGO, TODATE, and PERIODROLLING.

Figure 9–11 shows the Time option in the Logical Dimension dialog.

Figure 9–11 Time Option in Logical Dimension Dialog



Setting Chronological Keys for Each Level

Designate a chronological key for every level of each dimension hierarchy. This key must meet the requirements of being sequential (the members have a natural order), cardinal (all members are spaced the same distance apart at a given level, such as day or month), and complete (no members missing).


The Oracle BI Server uses the chronological key to create mathematically correct time series predictions, such as Jan + 2 months = Mar. You should set a chronological key for every level (except for the Grand Total level) so that you can perform time series operations on all levels with good performance. This enables you to use an AGO, TODATE, or PERIODROLLING function for any level of any time dimension hierarchy, such as fiscal month ago, calendar year ago, and day ago.

Theoretically, time series functions operate correctly if only the bottom level key in the Logical Dimension is chronological. In practice, however, this causes performance problems because it forces the physical query to use the lowest grain, causing joins of orders of magnitude more rows (for example, 365 times more rows for a "year ago" joining at the "day" grain). It also means higher-level aggregate tables are never selected by the query planner when using the time series functions, which again significantly slows the query.

As with any level key, be sure the key is unique at its level. For example, a column containing simple month names such as "January" is not unique unless it is concatenated to a column containing year names.

Figure 9–12 shows how to designate a chronological key in the Logical Level dialog.

Figure 9–12 Designating a Chronological Key in the Logical Level Dialog



Creating AGO, TODATE, and PERIODROLLING Measures

You can build time series measures by creating derived expressions from base measures. To do this, create a new logical column and select **Derived from existing columns using an expression**, then open Expression Builder to build the appropriate time series function.

Follow these guidelines when modeling time series functions:

- Time series functions cannot be derived from measures that use the fragmentation form of federation. This rule prevents some complex boundary conditions and cross-source assumptions in the query generation and result merging, such as the need to join some time dimension rows from one source to some of the fact rows in a different source.
- To reduce maintenance and increase accuracy, it is best to create a single base measure, and then derive a family of time series measures from it. For example, start with a base measure, then define variations for month-ago, year-ago, month-to-date, and so on. To do this, select **Derived from existing columns using an expression** and refer to the base measure in the expression.

Example 9–3 shows how to build the AGO measure. See Appendix C, "Logical SQL Reference" for detailed syntax for the other time series functions, TODATE and PERIODROLLING.


Example 9–3 Creating the AGO Measure

This example explains how to create one of the derived AGO measures in the Sampleapp demonstration repository.

1. In the Business Model and Mapping layer, create a new logical column. Name the column **2-04 Billed Qty (Mago)**.
2. In the Column Source tab, select **Derived from existing columns using an expression** and click the Expression Builder button.
3. In Expression Builder, select the Ago function to create a template for the arguments. To do this, select **Functions** in the Category pane, **Time Series Functions** in the Functions pane, and **Ago** in the Time Series Functions pane.

Figure 9–13 shows the AGO function in Expression Builder.

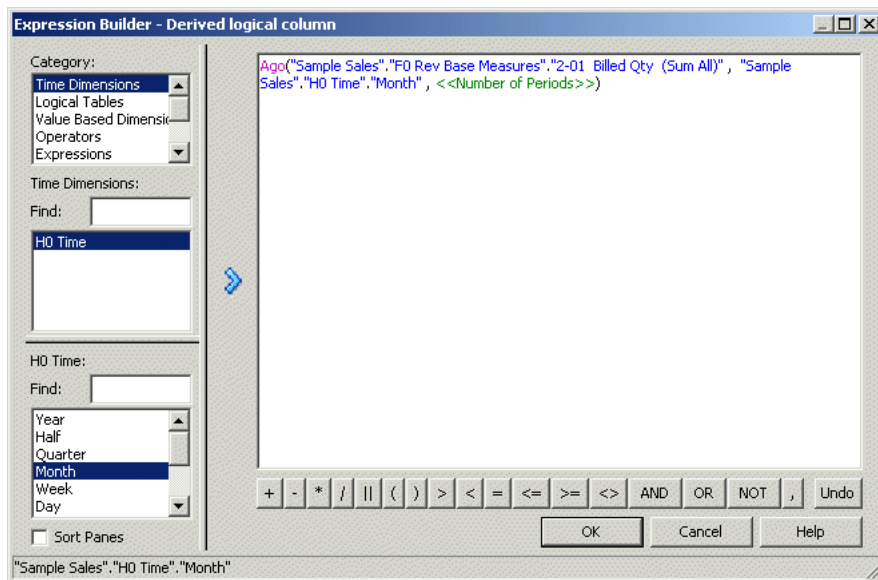
Figure 9–13 AGO Function in Expression Builder



4. Select the first argument, **Measure**, then use the selection panes to select the base measure from which to derive this column. In this example, select "Sample Sales"."F0 Rev Base Measures"."2-01 Billed Qty (Sum All)."
5. Select the second argument, **Level**, then use the selection panes to select the unit of the ago offset. It must be defined as a level of the time dimension, so that it can take advantage of the chronological keys built in the time dimension. In this example, select **Time Dimensions** in the Category pane, **HO Time** in the Time Dimensions pane, and **Month** in the HO Time pane.

Figure 9–14 shows the Month level in Expression Builder.

Figure 9–14 Selecting the Level Argument in Expression Builder



6. Select the third argument, Number of Periods, and enter the size of the offset you want to use for this measure. In this example, type 1.

7. Click **OK** in the Expression Builder dialog, then click **OK** in the Logical Column dialog.

Managing Logical Table Sources (Mappings)

Logical table sources define the mappings from a single logical table to one or more physical tables. The physical to logical mapping can also be used to specify transformations that occur between the Physical layer and the Business Model and Mapping layer, as well as to enable aggregate navigation and fragmentation.

One logical table source folder exists for each logical table. The folder contains one or more logical table sources. You can also view logical table sources from the Sources tab of the Logical Table dialog.

Logical tables can have many physical table sources. A single logical column might map to many physical columns from multiple physical tables, including aggregate tables that map to the column if a query asks for the appropriate level of aggregation on that column.

This chapter contains the following topics:

- [Creating Logical Table Sources](#)
- [Defining Physical to Logical Table Source Mappings and Creating Calculated Items](#)
- [Defining Content of Logical Table Sources](#)
- [Working with Parent-Child Settings in the Logical Table Source](#)
- [Setting Up Aggregate Navigation by Creating Sources for Aggregated Fact Data](#)
- [Setting Up Fragmentation Content for Aggregate Navigation](#)

Creating Logical Table Sources

When you create logical tables and columns by dragging and dropping from the Physical layer, the logical table sources are generated automatically. If you create the logical tables manually, you need to also create the sources manually.

You also add new logical table sources when multiple physical tables can be the source of information. For example, many tables could hold information for revenue. You might have three different business units (each with its own order system) where you get revenue information. In another example, you might periodically summarize revenue from an orders system or a financial system and use this table for high-level reporting.

Use the General tab of the Logical Table Source dialog to define general properties for the logical table source.

To create a logical table source:

1. In the Business Model and Mapping layer of the Administration Tool, right-click a logical table and select **New Object**, then select **Logical Table Source**.
2. In the Logical Table Source dialog, in the General tab, type a name for the logical table source.
3. Click the **Add** button. In the Browse dialog, you can view joins and select tables for the logical table source. When there are two or more tables in a logical table source, all of the participating tables must have joins defined between them.
4. To view the joins, in the Browse dialog, select a table and click **View**. After reviewing the joins in the Physical Table dialog, click **Cancel**.
5. To add tables to the table source, select the desired tables in the **Name** list and click **Select**.
6. Optionally, in the **Priority Group** field, enter a priority group number for this logical table source. See "[Setting Priority Group Numbers for Logical Table Sources](#)" for more information.
7. In the Logical Table Source dialog, click the Column Mapping tab and complete the fields. See "[Defining Physical to Logical Table Source Mappings and Creating Calculated Items](#)" for instructions.
8. In the Logical Table dialog, click the Content tab and complete the fields. See "[Defining Content of Logical Table Sources](#)" for instructions.
9. Click **OK**.

Setting Priority Group Numbers for Logical Table Sources

You can set priority group numbers to determine which logical table source should be used for queries for which there is more than one logical table source that can satisfy the requested set of columns.

For example, you might have user queries that can be fulfilled by both a data warehouse and an OLTP source. Often, access to an operational system is "expensive," while access to a data warehouse is "cheap." In this situation, you can assign a higher priority to the data warehouse to ensure that all queries are fulfilled by the data warehouse if possible.

Note that the priority group of a given logical table source does not always ensure that a particular query will be fulfilled by that source. Priority group assignments are only one of many factors used by the Oracle BI Server to determine which logical table source to select for a given query. However, the logical table source priority is the most significant metric and is considered before any other cost metric.

To assign priority group numbers, rank your logical table sources in numeric order, with 0 being the highest-priority source. You can assign the same number to multiple sources. For example, you can have two logical table sources in priority group 0, two logical table sources in priority group 1, and so on. Often, only two priority groups are necessary (0 and 1).

Assigning priority groups is optional. All logical table sources are set to priority 0 by default.

In some situations, you might want to allow users to reverse the normal logical table source priority ranking at query time. To accomplish this, you can use a combination of session variables and request variables with logical table source priority groups. This feature provides a way to dynamically select a source at run time, depending on user preference.

To enable this dynamic selection, you must first create the `REVERSIBLE_LTS_PRIORITY_SA_VEC` session variable in the repository. Create this variable as a string vector session variable that uses a row-wise session initialization block. `REVERSIBLE_LTS_PRIORITY_SA_VEC` should list the subject areas for which you want to allow users to reverse the logical table source priority ranking. You must define this variable to enable priority ranking reversal.

After you have defined the set of subject areas for which you want to allow priority ranking reversal, users can include the request variable `REVERSE_LTS_PRIORITY` with their queries to reverse the logical table source priority ranking. This request variable can be set to 1 to reverse the logical table source priority, or 0 to keep the normal logical table source priority.

As an alternative to using a request variable at query time, you can also define a predetermined set of subject areas for which the logical table source priority should be permanently reversed. To do this, create the session variable `REVERSED_LTS_PRIORITY_SA_VEC` in the repository. Create this variable as a string vector session variable that uses a row-wise session initialization block. `REVERSED_LTS_PRIORITY_SA_VEC` should list the subject areas for which you want the logical table source priority to be permanently reversed.

See "[Creating Session Variables](#)" for more information about how to define session variables in the Administration Tool.

Example of `REVERSIBLE_LTS_PRIORITY_SA_VEC`

You could create a table called `SA_TABLE` that contains two columns: `SUBJECT_AREA_NAME` and `REVERSIBLE`. This table could contain rows mapping subject area names to their reversible values (1 or 0), as follows:

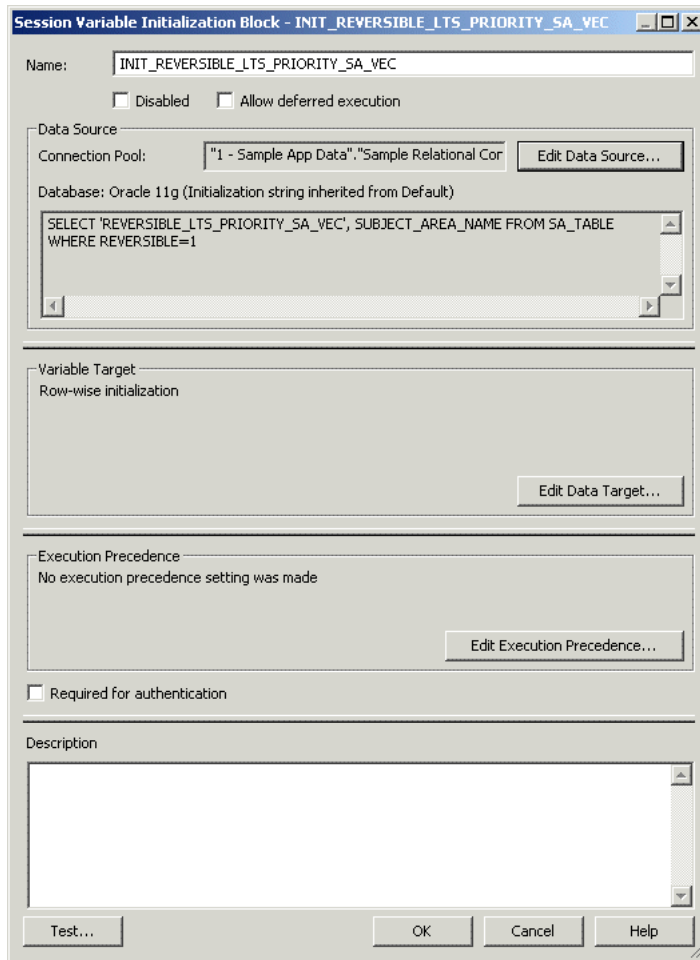
<code>SUBJECT_AREA_NAME</code>	<code>REVERSIBLE</code>
my_sa_1	1
my_sa_2	0

Then, you would create a string vector session variable called `REVERSIBLE_LTS_PRIORITY_SA_VEC` with a row-wise session initialization block. The initialization string for this initialization block could be similar to the following:

```
SELECT 'REVERSIBLE_LTS_PRIORITY_SA_VEC', SUBJECT_AREA_NAME FROM SA_TABLE
WHERE REVERSIBLE=1
```

[Figure 10–1](#) shows the Session Variable Initialization Block dialog for this example.

Figure 10–1 Session Variable Initialization Block Dialog for REVERSIBLE_LTS_PRIORITY_SA_VEC Example



Defining Physical to Logical Table Source Mappings and Creating Calculated Items

Use the Column Mapping tab of the Logical Table Source dialog to map logical columns to physical columns. The physical to logical mapping can also be used to specify transformations that occur between the Physical layer and the Business Model and Mapping layer. The transformations can be simple, such as changing an integer data type to a character, or more complex, such as applying a formula to find a percentage of sales per unit of population. Applying these transformations is typically referred to as creating calculated items.

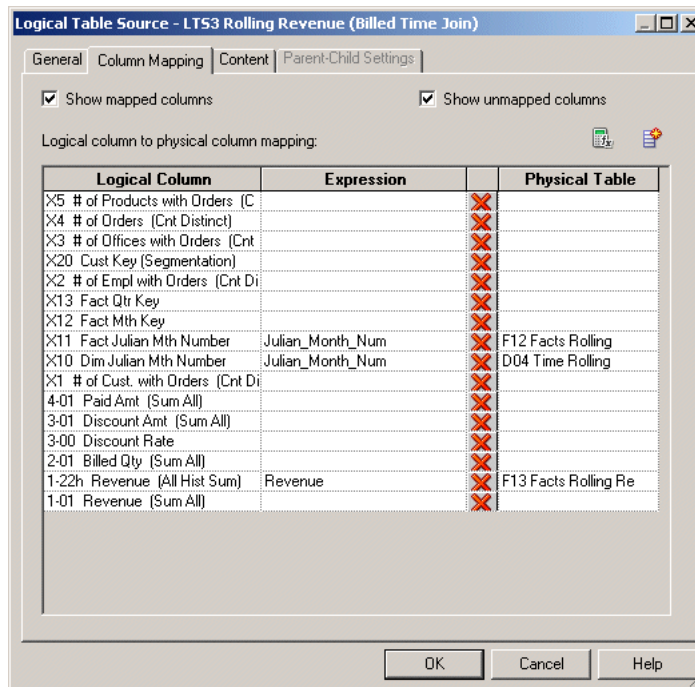
The data type of a logical column is determined by its logical table source mappings. For example, if a logical column has one physical source with a data type of VARCHAR(50) not-nullable, and another physical source with a data type of VARCHAR(20), nullable, then the data type of the logical column is VARCHAR(50) nullable. This final type is called a promoted type. Because of the rules governing logical table source mappings, you cannot map physical sources with data types that cannot be promoted (such as an INT with a VARCHAR).

To map logical columns to physical columns:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a logical table source.
2. In the Logical Table Source dialog, click the Column Mapping tab.
3. In the Column Mapping tab, maximize or enlarge the dialog to show all the contents, as shown in [Figure 10–2](#).

In the Column Mapping tab, in the **Logical column to physical column mapping** area, you can sort the rows (toggle among ascending order, descending order, and then restore original order) by clicking a column heading.

Figure 10–2 Column Mapping Tab of Logical Table Source Dialog



4. In the **Physical Table** column, select the table that contains the column you want to map.

When you select a cell in the Physical Table column, a list appears. It contains a list of tables currently included in this logical table source.

5. In the **Expression** column, select the physical column corresponding to each logical column.

When you select a cell in the **Expression** column, a list appears. It contains a list of physical columns currently included in this logical table source.

6. To open Expression Builder, click the **Expression Builder** button.

All columns used in creating physical expressions must be in tables included in the logical table source. You cannot create expressions involving columns in tables outside the source.

You can use Expression Builder to create calculated items, in which formulas are applied pre-aggregation. For example, to create the measure "tons sold" using the columns `units_sold` and `unit_weight`, you apply a pre-aggregation formula (`fact.units_sold*product.unit_weight`), and then apply the aggregation rule `SUM` in

the measure object. Another example is using `CAST` to transform a column of type `TIMESTAMP` to type `DATE` for faster display in Answers and other clients (for example, `CAST ("DB" . " . "TABLE" . "COL" AS DATE)`).

You can also conform sources by creating expressions that perform transformations on physical data. For example, you can use the `CAST` function to transform a column with a character data type to an integer data type, to match data coming from a second logical table source. Other examples include using `CONCATENATE` or math functions to make similar transformations on physical data.

See "[Creating Derived Columns](#)" for calculations that need to occur post-aggregation.

7. To remove a column mapping, click the **Delete** button. You might need to scroll to the right to locate the **Delete** button.
8. After you map the appropriate columns, click **OK**.

Unmapping a Logical Column from Its Source

In the Logical Column dialog, the Column Source tab contains information about the logical column. You can edit the logical table sources from which the column derives its data, or unmap it from its sources.

To unmap a logical column from its source:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a logical column.
2. In the Logical Column dialog, click the Column Source tab.
3. In the **Logical Table Source** list, select a source and click **Unmap**.
4. Click **OK**.

Defining Content of Logical Table Sources

To use a source correctly, the Oracle BI Server has to know what each source contains in terms of the business model. Therefore, you need to define aggregation content for each logical table source of a fact table. The aggregation content rule defines at what level of granularity the data is stored in this fact table. For each dimension that relates to this fact logical table, define the level of granularity, making sure that every related dimension is defined. See "[Setting Up Aggregate Navigation by Creating Sources for Aggregated Fact Data](#)" for more information.

If a logical table is sourced from a set of fragments, it is not required that every individual fragment maps the same set of columns. However, the server returns different answers depending on how columns are mapped.

- If all the fragments of a logical table map the same set of columns, than the set of fragmented sources is considered to be the whole universe of logical table sources for the logical table. This means that measure aggregations can be calculated based on the set of fragments.
- If the set of mapped columns differ across the fragments, than the Oracle BI Server assumes that it does not have the whole universe of fragments, and therefore it would be incorrect to calculate aggregate rollups (since some fragments are missing). In this case, the server returns `NULL` as measure aggregates.

Note: It is recommended that all the fragments map the same set of columns.

Use the Content tab of the Logical Table Source dialog to define any aggregate table content definitions, fragmented table definitions for the source, and WHERE clauses (if you want to limit the number of rows returned). See "[Setting Up Fragmentation Content for Aggregate Navigation](#)" for additional information.

Verifying that Joins Exist from Dimension Tables to Fact Table

This source content information tells the Oracle BI Server what it needs to know to send queries to the appropriate physical aggregate fact tables, joined to and constrained by values in the appropriate physical aggregate dimension tables. Be sure that joins exist between the aggregate fact tables and the aggregate dimension tables in the Physical layer.

One recommended way to verify joins is to select a fact logical table and open a Business Model Diagram (Selected Tables and Direct Joins). Only the dimension logical tables that are directly joined to this fact logical table appear in the diagram. It does not show dimension tables if the same physical table is used in logical fact and dimension sources.

Figure 10-3 shows an example of how the Fact - Assess fact logical table appears in a Business Model Diagram (Selected Tables and Direct Joins) view.

Figure 10-3 Business Model Diagram of Direct Joins for a Fact Table

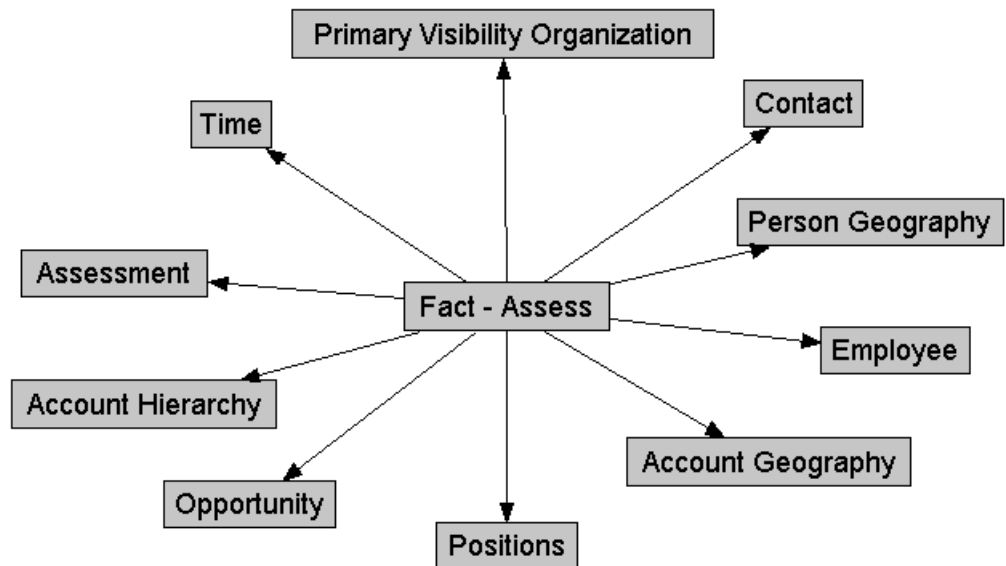


Table 10-1 contains a list of the logical level for each dimension table that is directly joined the Fact - Assess fact table shown in Figure 10-3.

Table 10-1 Dimension and Logical Level as Shown in Content Tab

Dimension	Logical Level
Account Geography	Postal Code Detail
Person Geography	Postal Code Detail

Table 10–1 (Cont.) Dimension and Logical Level as Shown in Content Tab

Dimension	Logical Level
Time	Day Detail
Account Organization	Account Detail
Opportunity	Oppty Detail
Primary Visibility Organization	Detail
Employee	Detail
Assessment	Detail
Contact (W_PERSON_D)	Detail
FINS Time	Day
Positions	Details

To create logical table source content definitions:

1. In the Business Model and Mapping layer of the Administration Tool, double-click a logical table source.
2. In the Logical Table Source dialog, click the Content tab and perform the following steps using [Table 10–2](#) as a guide.
3. If a logical source is an aggregate table and you have defined logical dimensions, select **Logical Level** from the **Aggregation content, group-by** list. Then, in the **Logical Level** list, select the appropriate level for each logical dimension table to which the logical fact table is joined.

You should specify a logical level for each dimension, unless you are specifying the Grand Total level. Dimensions with no level specified are interpreted as being at the most detailed level.

Caution: Although you have the option to specify aggregate content by logical level or column, it is recommended that you use logical levels exclusively. If you must define content by columns, do the following:

1. Select **Column** from the **Aggregation content, group-by** list.
2. In the **Table** pane, select each logical dimension table that defines the aggregation level of the source.
3. In the **Column** pane, select the logical column for each dimension that defines how the aggregations were grouped.

When there are multiple logical columns that could be used, select the one that maps to the key of the source physical table. For example, if data has been aggregated to the Region logical level, pick the logical column that maps to the key of the Region table.

Do not mix aggregation by logical level and column in the same business model.

4. To specify fragmented table definitions for the source, use the **Fragmentation content** box to describe the range of values included in the source when a source represents a portion of the data at a given level of aggregation.

You can type the formula directly into the box, or click the Expression Builder button to the right of the box. In the Expression Builder for Fragmentation Content, you can specify content in terms of existing logical columns. See ["Setting Up Fragmentation Content for Aggregate Navigation"](#) for additional information.

5. Select **This source should be combined with other sources at this level.**

This option is only for multiple sources that are at the same level of aggregation. For example, one logical table source might point to records for people with last names A-M, while a second logical table source might point to records for people with last names N-Z.

6. (Optional) To limit the number of rows the source uses in the resultant table, specify *WHERE* clause filters in the box labeled **Use this "WHERE clause" filter to limit rows returned (exclude the "WHERE")**. You can enter the *WHERE* clause directly, or you can click the Expression Builder button to open the Expression Builder, create the *WHERE* clause, and click **OK**.

See ["About WHERE Clause Filters"](#) for more information.

7. If the values for the source are unique, select the option **Select distinct values.**

Table 10–2 Content Tab Options for Logical Table Source

Options	Description
Aggregation content, group by	How the content is aggregated.
More button	<p>When you click More, the following options appear:</p> <ul style="list-style-type: none"> ▪ Copy. (Available only for fact tables) Copies aggregation content to the Windows clipboard. You can paste the Dimension.Level info into a text editor and use it for searching or for adding to documentation. Note that Copy is not available if the expression is empty. ▪ Copy from. (Available for fact tables and dimension tables) Copies aggregation content from another logical table source in the same business model. You need to specify the source from which to copy the aggregation content. (Multiple business models appear but only the logical table sources from the current business model are selectable.) ▪ Get Levels. (Available only for fact tables) Changes aggregation content. If joins do not exist between fact table sources and dimension table sources (for example, if the same physical table is in both sources), the aggregation content determined by the Administration Tool does not include the aggregation content of this dimension. ▪ Check Levels. (Available only for fact tables) Checks the aggregation content of logical fact table sources (not dimension table sources). The information returned depends on the existence of dimensions and hierarchies with logical levels and level keys, and physical joins between tables in dimension table sources and the tables in the fact table source. (If the same tables exist in the fact and dimension sources and there are no physical joins between tables in the sources, Check Levels does not include the aggregation content of this dimension.)

Table 10–2 (Cont.) Content Tab Options for Logical Table Source

Options	Description
Fragmentation content	A description of the contents of a data source in business model terms. Data is fragmented when information at the same level of aggregation is split into multiple tables depending on the values of the data. A common situation would be to have data fragmented by time period. See " Setting Up Fragmentation Content for Aggregate Navigation " for additional information.
This source should be combined with other sources at this level	Select this option when data sources at the same level of aggregation do not contain overlapping information. In this situation, all sources must be combined to get a complete picture of information at this level of aggregation.
Select distinct values	Used if the values for the source are unique.

About WHERE Clause Filters

The `WHERE` clause filter is used to constrain the physical tables referenced in the logical table source. If there are no constraints on the aggregate source, leave the `WHERE` clause filter blank.

Each logical table source should contain data at a single intersection of aggregation levels. You would not want to create a source, for example, that had sales data at both the Brand and Manufacturer levels. If the physical tables include data at multiple levels, add an appropriate `WHERE` clause constraint to filter values to a single level.

Any constraints in the `WHERE` clause filter are made on the physical tables in the source.

Working with Parent-Child Settings in the Logical Table Source

Sometimes, a logical table is part of a dimension with a parent-child hierarchy that is based on relational tables. When this is the case, the logical table includes both a physical source and a source for the parent-child relationship table required for the parent-child hierarchy. Parent-child relationship tables explicitly define the inter-member relationships for parent-child hierarchies.

Typically, logical table sources for parent-child relationship tables are created automatically when you run the scripts created by the Generate Parent-Child Table Wizard. You access this wizard from the Parent-Child Table Settings dialog, available in the dimension object.

Note: The Generate Parent-Child Table Wizard feature is not available from the Logical Table Source dialog. You must go to the dimension object to create scripts to generate the parent-child relationship table.

You can view details for the parent-child relationship table source in the Parent-Child Settings tab of the Logical Table Source dialog. The following information appears in the tab:

- **Parent-Child Table:** Shows the name of the parent-child relationship table on which this source is based.
- **Member Key:** The name of the column in the parent-child relationship table that identifies the member.

- **Parent Key:** The name of the column in the parent-child relationship table that identifies an ancestor of the member.
- **Relationship Distance:** The name of the column in the parent-child relationship table that specifies the number of parent-child hierarchical levels from the member to the ancestor.
- **Leaf Node Identifier:** The name of the column in the parent-child relationship table that indicates if the member is a leaf member (1=Yes, 0=No).

See "[Creating Dimensions with Parent-Child Hierarchies](#)" for more information about parent-child relationship tables.

Setting Up Aggregate Navigation by Creating Sources for Aggregated Fact Data

Aggregate tables store precomputed results from measures that have been aggregated over a set of dimensional attributes. Each aggregate table column contains data at a given set of levels. For example, a monthly sales table might contain a precomputed sum of the revenue for each product in each store during each month. You configure this metadata in the Content tab of the Logical Table Source dialog.

When you create a logical table source for an aggregate fact table, you should create corresponding logical dimension table sources at the same levels of aggregation.

You need to have at least one logical dimension table source for each level of aggregation. If the sources at each level already exist, you do not need to create new ones.

For example, you might have a monthly sales fact table containing a precomputed sum of the revenue for each product in each store during each month. You need to have the following three other dimension sources, one for each of the logical dimension tables referenced in the example:

- A source for the Product logical table with one of the following content specifications:
 - By logical level: ProductDimension.ProductLevel
 - By column: Product.Product_Name
- A source for the Store logical table with one of the following content specifications:
 - By logical level: StoreDimension.StoreLevel
 - By column: Store.Store_Name
- A source for the Time logical table with one of the following content specifications:
 - By logical level: TimeDimension.MonthLevel
 - By column: Time.Month

At query time, the Oracle BI Server first determines which sources have enough detail to answer the query. Out of these sources, the Oracle BI Server chooses the most aggregated source to answer the query, because it is assumed to be the fastest. The most aggregated source is the one with the lowest multiplied number of elements. See "[Creating Logical Levels in a Dimension](#)" for more information about specifying the number of elements at each level.

Setting Up Fragmentation Content for Aggregate Navigation

When a logical table source does not contain the entire set of data at a given level, you need to specify the portion, or fragment, of the set that it does contain. You describe the content in terms of logical columns in the **Fragmentation content** box in the Content tab of the Logical Table Source dialog.

The examples in this section illustrate techniques and rules for specifying the fragmentation content of sources.

This section contains the following topics:

- [Specifying Fragmentation Content for Single Column, Value-Based Predicates](#)
- [Specifying Fragmentation Content for Single Column, Range-Based Predicates](#)
- [Specifying Fragmentation Content for Aggregate Table Fragments](#)

Specifying Fragmentation Content for Single Column, Value-Based Predicates

The IN predicates can be replaced with either an equality predicate or multiple equality predicates separated by the OR connective.

Fragment 1:

```
logicalColumn IN <valueList1>
```

Fragment n:

```
logicalColumn IN <valueListn>
```

Specifying Fragmentation Content for Single Column, Range-Based Predicates

Fragment 1:

```
logicalColumn >= valueof(START_VALUE) AND logicalColumn < valueof(MID_VALUE1)
```

Fragment 2:

```
logicalColumn >= valueof(MID_VALUE1) AND logicalColumn < valueof(MID_VALUE2)
```

Fragment n:

```
logicalColumn >= valueof(MID_VALUEn-1) AND logicalColumn < valueof(END_VALUE)
```

Pick your start point, midpoints, and endpoint carefully.

Note: Use >= and < predicates to make sure the fragment content descriptions do not overlap. For each fragment, the upper value must be expressed as <. You will get an error if you use <=. Likewise, you cannot use the BETWEEN predicate to describe fragment range content.

The valueof referenced here is the value of a repository variable. If you use repository values in your expression, note that the following construct does not work for Fragment 2:

```
logicalColumn >= valueof(MID_VALUE1)+1 AND logicalColumn < valueof(MID_VALUE2)
```

Use another repository variable instead of valueof(MID_VALUE₁)+1.

The same variables, for example, `valueof (MID_VALUE1)`, do not have to appear in the content of both fragments. You could set another variable, and create statements of the following form:

Fragment 1:

```
logicalColumn >= valueof (START_VALUE) AND logicalColumn < valueof (MID_VALUE1)
```

Fragment 2:

```
logicalColumn >= valueof (MID_VALUE2) AND logicalColumn < valueof (MID_VALUE3)
```

For more information about variables, see [Chapter 18](#).

Specifying Multicolumn Content Descriptions

An arbitrary number of predicates on different columns can be included in each content filter. Each column predicate can be value-based or range-based.

Fragment 1:

```
<logicalColumn1 predicate> AND <logicalColumn2 predicate > ... AND <logicalColumnM predicate>
```

Fragment n:

```
<logicalColumn1 predicate> AND <logicalColumn2 predicate > ... AND <logicalColumnM predicate>
```

Ideally, all fragments will have predicates on the same M columns. If there is no predicate constraint on a logical column, the Oracle BI Server assumes that the fragment contains data for all values in that logical column. See "[Specifying Parallel Content Descriptions](#)" for exceptions using the OR predicate.

Specifying Parallel Content Descriptions

Unfortunately, the preceding techniques are still not sufficient to handle dates because of the multiple hierarchical relationships across logical columns, such as year > year month > date; month > year month > date. For example, consider fragments delineated by different points in time, such as year and month. Constraining sufficiently far back on year should be enough to drive the selection of just the historical fragment. The parallel OR technique supports this, as shown in the next example. This example assumes that the snapshot month was April 1, 12:00 a.m. in the year 1999.

Fragment 1 (Historical):

```
EnterpriseModel.Period."Day" < VALUEOF("Snapshot Date") OR
EnterpriseModel.Period.MonthCode < VALUEOF("Snapshot Year Month") OR
EnterpriseModel.Period."Year" < VALUEOF("Snapshot Year") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
    EnterpriseModel.Period."Month in Year" < VALUEOF("Snapshot Month") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
    EnterpriseModel.Period."Monthname" IN ('Mar', 'Feb', 'Jan')
```

Fragment 2 (Current):

```
EnterpriseModel.Period."Day" >= VALUEOF("Snapshot Date") OR
EnterpriseModel.Period.MonthCode >= VALUEOF("Snapshot Year Month") OR
EnterpriseModel.Period."Year" > VALUEOF("Snapshot Year") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
    EnterpriseModel.Period."Month in Year" >= VALUEOF("Snapshot Month") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
```

```
EnterpriseModel.Period."Monthname" IN ('Dec', 'Nov', 'Oct', 'Sep', 'Aug', 'Jul',
'Jun', '', 'Apr')
```

If the logical model does not go down to the date level of detail, then omit the predicate on `EnterpriseModel.Period."Day"` in the preceding example.

Notice the use of the `OR` connective to support parallel content description tracks.

Examples of Parallel Content Descriptions In this section, the Track *n* labels in the examples are shown to help relate the examples to the discussion that follows. You would not include these labels in the actual fragmentation content statement.

Example 10–1 Fragment 1 (Historical)

```
Track 1 EnterpriseModel.Period."Day" < VALUEOF("Snapshot Date") OR
Track 2 EnterpriseModel.Period.MonthCode < VALUEOF("Snapshot Year Month") OR
Track 3 EnterpriseModel.Period."Year" < VALUEOF("Snapshot Year") OR
Track 4 EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
      EnterpriseModel.Period."Month in Year" < VALUEOF("Snapshot Month") OR
Track 5 EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
      EnterpriseModel.Period."Monthname" IN ('Mar', 'Feb', 'Jan')
```

For example, consider the first track on `EnterpriseModel.Period."Day."` In the historical fragment, the `<` predicate tells the Oracle BI Server that any queries that constrain on `Day` before the `Snapshot Date` fall within the historical fragment. Conversely, the `>=` predicate in the current fragment on `Day` indicates that the current fragment does not contain data before the `Snapshot Date`.

The second track on `MonthCode` (for example, 199912) is similar to `Day`. It uses the `<` and `>=` predicates, as there is a nonoverlapping delineation on month (because the snapshot date is April 1). The key rule to remember is that each additional parallel track must reference a different column set. Common columns can be used, but the overall column set must be unique. The Oracle BI Server uses the column set to select the most appropriate track.

The third track on `Year` (`<` in the historical fragment and `>` in the current fragment) tells the Oracle BI Server that optimal (single) fragment selections can be made on queries that just constrain on year. For example, a logical query on `Year IN (1997, 1998)` should only hit the historical fragment. Likewise, a query on `Year = 2000` should only hit the current fragment. However, a query that hits the year 1999 cannot be answered by the content described in this track, and therefore hits both fragments, unless additional information can be found in subsequent tracks.

The fourth track describes the fragment set for `Year` and `Month in Year` (month integer). Notice the use of the multicolumn content description technique, described previously. Notice the use of `<` and `>=` predicates, as there is no ambiguity or overlap for these two columns.

The fifth track describes fragment content in terms of `Year` and `Monthname`. It uses the value-based `IN` predicate technique.

As an embellishment, suppose the snapshot date fell on a specific day within a month: therefore, multicolumn content descriptions on just year and month would overlap on the specific snapshot month. To specify this ambiguity, `<=` and `>=` predicates are used.

Fragment 1 (Historical):

```
EnterpriseModel.Period."Day" < VALUEOF("Snapshot Date") OR
EnterpriseModel.Period.MonthCode <= VALUEOF("Snapshot Year Month") OR
EnterpriseModel.Period."Year" < VALUEOF("Snapshot Year") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
```

```
EnterpriseModel.Period."Month in Year" <= VALUEOF("Snapshot Month") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
EnterpriseModel.Period."Monthname" IN ('Apr', 'Mar', 'Feb', 'Jan')
```

Fragment 2 (Current):

```
EnterpriseModel.Period."Day" >= VALUEOF("Snapshot Date") OR
EnterpriseModel.Period.MonthCode >= VALUEOF("Snapshot Year Month") OR
EnterpriseModel.Period."Year" > VALUEOF("Snapshot Year") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
EnterpriseModel.Period."Month in Year" >= VALUEOF("Snapshot Month") OR
EnterpriseModel.Period."Year" = VALUEOF("Snapshot Year") AND
EnterpriseModel.Period."Monthname" IN ('Dec', 'Nov', 'Oct', 'Sep', 'Aug', 'Jul',
'Jun', '', 'Apr')
```

Specifying Unbalanced Parallel Content Descriptions

In an order entry application, time-based fragmentation between historical and current fragments is typically insufficient. For example, records might still be volatile, even though they are historical records entered into the database before the snapshot date.

Assume, in the following example, that open orders can be directly updated by the application until the order is shipped or canceled. After the order has shipped, however, the only change that can be made to the order is to type a separate compensating return order transaction.

There are two parallel tracks in the following content descriptions. The first track uses the multicolumn, parallel track techniques described in the preceding section. Notice the parentheses nesting the parallel calendar descriptions within the Shipped-or-Canceled order status multicolumn content description.

The second parallel track is present only in the Current fragment and specifies that all Open records are in the Current fragment only.

Fragment 1 (Historical):

```
Marketing."Order Status"."Order Status" IN ('Shipped', 'Canceled') AND
Marketing.Calendar."Calendar Date" <= VALUEOF("Snapshot Date") OR
Marketing.Calendar."Year" <= VALUEOF("Snapshot Year") OR
Marketing.Calendar."Year Month" <= VALUEOF("Snapshot Year Month")
```

Fragment 2 (Current):

```
Marketing."Order Status"."Order Status" IN ('Shipped', 'Canceled') AND
Marketing.Calendar."Calendar Date" > VALUEOF("Snapshot Date") OR
Marketing.Calendar."Year" >= VALUEOF("Snapshot Year") OR
Marketing.Calendar."Year Month" >= VALUEOF("Snapshot Year Month") OR
Marketing."Order Status"."Order Status" = 'Open'
```

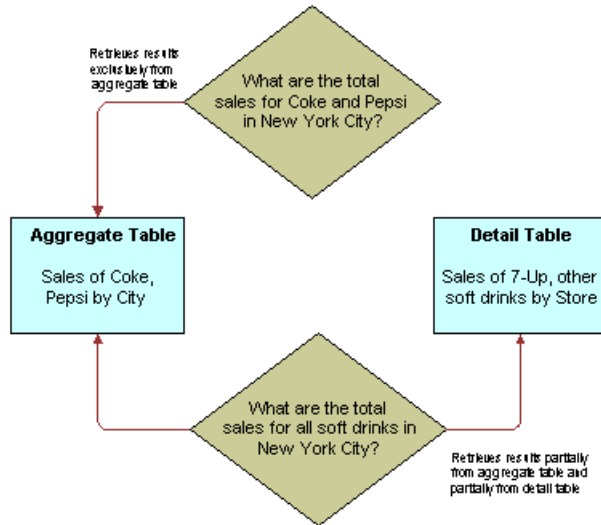
The overlapping Year and Month descriptions in the two fragments do not cause a problem, as overlap is permissible when there are parallel tracks. The rule is that at least one of the tracks has to be nonoverlapping. The other tracks can have overlap.

Specifying Fragmentation Content for Aggregate Table Fragments

Information at a given level of aggregation is sometimes stored in multiple physical tables. When individual sources at a given level contain information for a portion or fragment of the domain, the Oracle BI Server needs to know the content of the sources in order to pick the appropriate source for the query.

For example, suppose you have a database that tracks the sales of soft drinks in all stores. The detail level of data is at the store level. Aggregate information, as described in Figure 10-4, is stored at the city level for the sales of Coke and Pepsi, but there is no aggregate information for the sales of 7-Up or any other of the sodas.

Figure 10-4 Aggregating Information



The goal of this type of configuration is to maximize the use of the aggregate table. If a query asks for sales figures for Coke and Pepsi, the data should be returned from the aggregate table. If a query asks for sales figures for all soft drinks, the aggregate table should be used for Coke and Pepsi and the detail data for the other brands.

The Oracle BI Server handles this type of partial aggregate navigation. To configure a repository to use aggregate fragments for queries whose domain spans multiple fragments, you need to define the entire domain for each level of aggregate data, even if you must configure an aggregate fragment as being based on a less summarized physical source.

This section contains the following topics:

- [Specifying the Aggregate Table Content](#)
- [Defining a Physical Layer Table with a Select Statement to Complete the Domain](#)
- [Specifying the SQL Virtual Table Content](#)
- [Creating Physical Joins for the Virtual Table](#)

Specifying the Aggregate Table Content

You configure the aggregate table navigation in the logical table source mappings. In the soft drink example, the aggregate table contains data for Coke and Pepsi sales at the city level. Its Aggregate content specification (in the Content tab of the Logical Table Source window) is similar to the following:

Group by logical level:

```
GeographyDim. CityLevel, ProductDim.ProductLevel
```

Its Fragmentation content specification (also in the Content tab of the Logical Table Source dialog) is similar to the following:

```
SoftDrinks.Products.Product IN ('Coke', 'Pepsi')
```

This content specification tells the Oracle BI Server that the source table has data at the city and product level for two of the products. Additionally, because this source is a fragment of the data at this level, you must select **This source should be combined with other sources at this level**, in the Content tab of the Logical Table Source dialog, to indicate that the source combines with other sources at the same level.

Defining a Physical Layer Table with a Select Statement to Complete the Domain

The data for the rest of the domain (the other types of sodas) is all stored at the store level. To define the entire domain at the aggregate level (city and product, in this example), you need to have a source that contains the rest of the domain at this level. Because the data at the store level is at a lower (that is, more detailed) level than at the city level, it is possible to calculate the city and product level detail from the store and product detail by adding up the product sales data of all of the stores in a city. This can be done in a query involving the store and product level table.

One way to do this is to define a table in the Physical layer with a Select statement that returns the store level calculations. To define the table, create a table in the Physical layer by right-clicking the physical schema object that the `SELECT` statement will be querying and selecting **New Physical Table**. Choose **Select** from the **Table Type** list, and type the SQL statement in the **Default Initialization String** box.

The SQL statement must define a virtual table that completes the domain at the level of the other aggregate tables. In this case, there is one existing aggregate table, and it contains data for Coke and Pepsi by city. Therefore, the SQL statement has to return all of the data at the city level, except for the Coke and Pepsi data.

Specifying the SQL Virtual Table Content

Next, create a new logical table source for the Sales column that covers the remainder of the domain at the city and product level. This source contains the virtual table created in the previous section. Map the Dollars logical column to the USDollars physical column in this virtual table.

The Aggregate content specification (in the Content tab of the Logical Table Source dialog) for this source is:

Group by logical level:

```
GeographyDim.CityLevel, ProductDim.ProductLevel
```

This tells the Oracle BI Server this source has data at the city and product level.

The Fragmentation content specification might be:

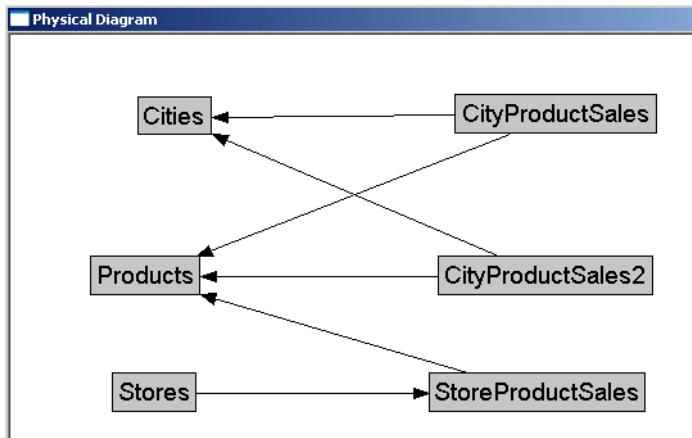
```
SoftDrinks.Products.Product = '7-Up'
```

Additionally, because it combines with the aggregate table containing the Coke and Pepsi data at the city and product level to complete the domain, you need to select the option in the Content tab of the Logical Table Source dialog indicating that the source is combined with other sources at the same level.

Creating Physical Joins for the Virtual Table

Construct the correct physical joins for the virtual table. Notice that CityProductSales2 joins to the Cities and Products tables in [Figure 10-5](#).

Figure 10–5 Example Physical Joins



In this example, the two sources comprise the whole domain for soda sales. A domain can have many sources. The sources have to all follow the rule that each level must contain sources that, when combined, comprise the whole domain of values at that level. Setting up the entire domain for each level helps ensure that queries asking for Coke, Pepsi, and 7-Up do not leave out 7-Up. It also helps ensure that queries requesting information that has been precomputed and stored in aggregate tables can retrieve that information from the aggregate tables, even if the query requests other information that is not stored in the aggregate tables.

Creating and Maintaining the Presentation Layer

The Presentation layer provides a way to present customized, secure, role-based views of a business model to users. Role-based views provide object security and also provide a way to hide some of the complexity of the business model. The Presentation layer also provides some of the functionality of the metadata model, such as the ability to set an implicit fact column.

Presentation layer views are called subject areas (formerly called presentation catalogs). You can have a subject area that is identical to your business model, or you can provide smaller, role-based subject areas that show a single subject, or that support a particular business role. Create subject areas that help you organize your content in a way that makes sense for your users.

After you have created the Business Model and Mapping layer, you can drag and drop entire business models to the Presentation layer in the Administration Tool. Alternatively, you can create subject areas and other Presentation layer objects manually.

This chapter explains how to use the Administration Tool to create and edit objects in the Presentation layer of a repository.

This chapter contains the following topics:

- [Creating and Customizing the Presentation Layer](#)
- [Working with Subject Areas](#)
- [Working with Presentation Tables and Columns](#)
- [Working with Presentation Hierarchies and Levels](#)
- [Setting Permissions for Presentation Layer Objects](#)
- [Creating Aliases \(Synonyms\) for Presentation Layer Objects](#)

Creating and Customizing the Presentation Layer

The Presentation layer provides a way to present customized views of a business model to users. Subject areas in the Presentation layer appear as catalogs to client tools that use the Oracle BI Server as an ODBC data source. Subject areas contain presentation tables, columns, hierarchies, and levels.

Even though the Logical SQL requests from Answers and other clients query the presentation tables and columns, the real logic for entities, relationships, joins, and so on is in the Business Model and Mapping layer. The primary function of the

Presentation layer is to provide custom names, dictionary entries, organization, and security for different groups of users.

This section contains the following topics:

- [Creating Subject Areas](#)
- [Removing Unneeded or Unwanted Columns](#)
- [Renaming Presentation Columns to User-Friendly Names](#)
- [Exporting Logical Keys in the Subject Area](#)
- [Setting an Implicit Fact Column in the Subject Area](#)
- [Maintaining the Presentation Layer](#)

Creating Subject Areas

There are several ways to create subject areas in the Presentation layer. The recommended method is to drag and drop a business model from the Business Model and Mapping layer to the Presentation layer, and then modify the Presentation layer based on what you want users to see. You can move columns between presentation tables, remove columns that do not need to be seen by the users, or even present all of the data in a single presentation table. You can create presentation tables to organize and categorize measures in a way that makes sense to your users.

You can also duplicate an existing subject area and its corresponding business model. See "[Duplicating a Business Model and Subject Area](#)" for more information.

Although each subject area must be populated with contents from a single business model, you can create multiple subject areas for one business model. For very large business models, you may want to do this to help users work with the content. Users in Oracle BI Answers can create queries that span multiple subject areas, as long as the subject areas correspond to the same business model.

There are many ways to create multiple subject areas from a single business model. One method is to drag a particular business model to the Presentation layer multiple times, then edit the properties or objects of the resulting subject areas as needed.

For example, if you have a business model called `ABC` that contains the `Geography` and `Products` dimensions, you can drag it to the Presentation layer twice. Two subject areas are created, with the default names `ABC` and `ABC#1`. You can then edit the subject areas as follows:

- Rename the `ABC` subject area to `DEF`, then delete the `Geography` presentation hierarchy
- Rename the `ABC#1` subject area to `XYZ`, then delete the `Products` presentation hierarchy

Users in Oracle BI Answers can then run queries that span both the `DEF` subject area (containing the `Products` hierarchy), and the `XYZ` subject area (containing the `Geography` hierarchy).

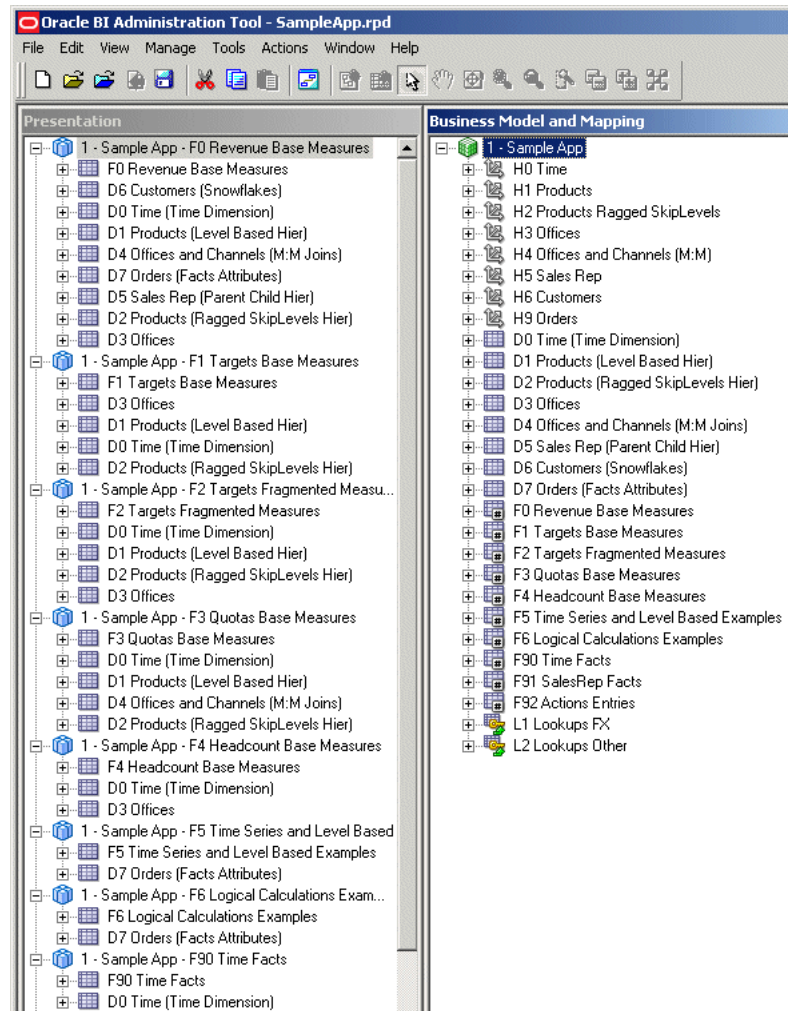
Automatically Creating Subject Areas Based on Logical Stars and Snowflakes

You can automatically create one subject area for each logical star or logical snowflake in your business model. Logical stars and logical snowflakes are both composed of a centralized fact table connected to multiple dimension tables. This feature provides another way to create multiple subject areas from a single business model.

To create a subject area for each fact table that is part of a logical star or snowflake, right-click the business model and select **Create Subject Areas for Logical Stars and Snowflakes**. The new subject areas are automatically created, each containing a fact table and only the dimension tables with which it is associated. This option is available for any business model that contains logical stars or logical snowflakes.

For example, if you choose this option for the SampleApp business model with nine fact tables, nine corresponding subject areas are created, each with one fact table and its associated dimension tables. Subject areas are also created for lookup tables. [Figure 11-1](#) shows exactly how the logical fact tables and dimension tables are modeled in the Presentation layer.

Figure 11-1 *Creating Subject Areas for Logical Stars and Snowflakes in SampleApp*



Removing Unneeded or Unwanted Columns

One important reason to use a custom Presentation layer is to make the schema as easy to use and understand as possible. Therefore, users should not be able to view columns that have no meaning to them. The following columns are examples of columns that you might want to remove from the Presentation layer:

- Key columns that have no business meaning

- Columns that users do not need to view (for example, codes, when text descriptions exist)
- Columns that users are not authorized to see

Renaming Presentation Columns to User-Friendly Names

By default, presentation columns have the same name as the corresponding logical column in the Business Model and Mapping layer. It is recommended to keep presentation column names and their source logical column names synchronized to reduce maintenance. To do this, ensure that **Use Logical Column Name** is selected in the Presentation Column dialog.

In some cases, however, you may want a different presentation column name to be shown to users. To do this, change the name of the presentation column in the Presentation Column dialog.

When you change the name of a presentation column, an alias is automatically created for the old name, so compatibility to the old name remains. See "[Creating Aliases \(Synonyms\) for Presentation Layer Objects](#)" for more information about aliases.

Note that you cannot rename a Presentation layer object to a name that is already in use as an alias for an object of the same type.

Exporting Logical Keys in the Subject Area

For each subject area in the Presentation layer, you can decide whether to export any logical keys as key columns to tools that access it. Exporting logical keys is irrelevant to users of Oracle BI Presentation Services, but it may be advantageous for some query and reporting tools. If you decide to export logical keys, make sure that the logical key columns exist in the table folders. In this situation, your business model should use logical key/foreign key joins.

When you select the option **Export logical keys** in the Subject Area dialog, any columns in the Presentation layer that are key columns in the Business Model and Mapping layer are listed as key columns to any ODBC client. This is the default selection. In most situations, this option should be selected.

Note: If you are using a tool that issues parameterized SQL queries, such as Microsoft Access, do not select the option **Export logical keys**. This stops the tool from issuing parameterized queries.

Setting an Implicit Fact Column in the Subject Area

For each subject area in the Presentation layer, you can set an implicit fact column. The implicit fact column is added to a query when it contains columns from two or more dimension tables and no measures. The column is not visible in the results. It is used to specify a default join path between dimension tables when there are several possible alternatives or contexts.

Maintaining the Presentation Layer

There is no automatic way to synchronize all changes between the Business Model and Mapping layer and the Presentation layer. For example, if you add logical columns to an existing logical table, or edit existing columns, you must manually update the corresponding Presentation layer objects.

However, the Administration Tool can automatically synchronize the name of presentation columns with their corresponding logical column names. To take advantage of this feature, ensure that **Use Logical Column Name** is selected in the Presentation Column dialog.

In some cases, if there are many changes to a logical table or even to an entire business model, it is easiest to delete the corresponding presentation table or subject area, and then and drag and drop the updated logical objects to the Presentation layer. For this reason, it is best to wait until the Business Model and Mapping layer is relatively stable before adding customizations in the Presentation layer.

Working with Subject Areas

In the Presentation layer, subject areas enable you to show different views of a business model to different sets of users. Subject areas have to be populated with contents from a single business model. They cannot span business models.

Typically, subject areas are created automatically by dragging and dropping business models from the logical layer.

To edit the properties of a subject area:

1. In the Presentation layer, double-click a subject area. The Subject Area dialog is displayed.
2. In the General tab, you can change the name for the subject area. Note that aliases are created automatically whenever presentation objects are renamed, so that any queries using the original name do not break.

Also, a subject area cannot have the same name as any of its child presentation tables. For example, you cannot have a subject area called Customer that has a Customer table within it.

3. To set permissions for this subject area, click **Permissions**. See "[Setting Permissions for Presentation Layer Objects](#)" for more information.
4. Select **Custom display name** or **Custom description** to dynamically display a custom name or custom description based on a session variable, typically for localization purposes. See "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization.
5. The **Business model** list displays the business model for this subject area.
6. To expose the logical keys to other applications, select the option **Export logical keys**.

In most situations, this option should be selected. Many client tools differentiate between key and nonkey columns, and the option **Export logical keys** provides client tools access to the key column metadata. Any join conditions the client tool adds to the query, however, are ignored, because the Oracle BI Server uses the joins defined in the repository.

Note: If you are using a tool that issues parameterized SQL queries, such as Microsoft Access, do not select the **Export logical keys** option. Not exporting logical keys stops the tool from issuing parameterized queries.

7. Optionally, you can set an **Implicit Fact Column**. This column is added to a query when it contains columns from two or more dimension tables and no measures. The column is not visible in the results. It is used to specify a default join path between dimension tables when there are several possible alternatives or contexts.
8. Optionally, type a description. This description appears in a mouse-over tooltip for the subject area in Oracle BI Answers.
9. In the Presentation Tables tab, you can add, remove, edit, or reorder the presentation tables for this subject area.
10. Use the Aliases tab to specify or delete aliases for this subject area. See "[Creating Aliases \(Synonyms\) for Presentation Layer Objects](#)" for more information about aliases.
11. Click OK.

Working with Presentation Tables and Columns

Presentation tables and presentation columns appear as folders and columns in Oracle BI Answers. You can customize presentation tables and presentation columns to help users craft queries based on their business needs.

This section contains the following topics:

- [Creating and Managing Presentation Tables](#)
- [Creating and Managing Presentation Columns](#)

Creating and Managing Presentation Tables

You can use presentation tables to organize columns into categories that make sense to the user community. A presentation table can contain columns from one or more logical tables. The names and object properties of the presentation tables are independent of the logical table properties.

Typically, presentation tables are created automatically by dragging and dropping logical tables from the logical layer.

To edit the properties of a presentation table:

1. In the Presentation layer, double-click a presentation table. The Presentation Table dialog appears.
2. In the General tab, you can change the name for the presentation table. Note that aliases are created automatically whenever presentation objects are renamed, so that any queries using the original name do not break.

Also, a presentation table cannot have the same name as its parent subject area. For example, you cannot have a subject area called Customer that has a Customer table within it.

3. To set permissions for this presentation table, click **Permissions**. See "[Setting Permissions for Presentation Layer Objects](#)" for more information.
4. Select **Custom display name** or **Custom description** to dynamically display a custom name or custom description based on a session variable, typically for localization purposes. See "Localizing Oracle Business Intelligence" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization.

5. In the Columns tab, you can add, remove, edit, or reorder the presentation columns for this presentation table.
6. In the Hierarchies tab, you can add, remove, edit, or reorder the presentation hierarchies for this presentation table.
7. Use the Aliases tab to specify or delete aliases for this presentation table. See ["Creating Aliases \(Synonyms\) for Presentation Layer Objects"](#) for more information about aliases.
8. Click **OK**.

To reorder a table or sort all tables in a subject area:

1. In the Presentation layer, double-click a subject area.
2. In the Subject Area dialog, click the Presentation Tables tab.
3. To move a table, in the **Name** list, select the table you want to reorder. Then, use drag-and-drop to reposition the table, or click the **Up** and **Down** buttons.
4. To sort all tables in alphanumeric order, click the **Name** column heading. This toggles the sort between ascending and descending alphanumeric order.

Nesting Folders in Answers

You can use the Administration Tool to update Presentation layer metadata to give the appearance of nested folders in Answers. To do this, open the Properties dialog for the presentation table that corresponds to the folder you want to nest and add -> to the beginning of the Description field.

For example, to nest the Sales Facts folder in the Facts folder in Answers, place the Sales Facts presentation table directly after the Facts presentation table in the metadata and add -> to the Description field in the Presentation Table - Sales Facts dialog. To nest a second folder called Marketing Facts in the Facts folder, add -> to the Description field for the Marketing Facts presentation table and place it directly after Sales Facts. Only two levels of nesting are supported.

Alternatively, you can prefix the name of the presentation table to be nested with a hyphen and a space and then place the table after the presentation table in which it nests. For example, to nest Sales Facts within Facts, place Sales Facts directly after Facts and change its name to - Sales Facts. When Answers displays the folder name in the left pane, it omits the hyphen and space from the folder name. However, the hyphen and space are visible to ODBC clients.

Note that these techniques do not provide true folder nesting. If you move the parent presentation table, the "child" presentation tables do not move with it. These techniques only provide the appearance of nesting in Answers.

Creating and Managing Presentation Columns

The presentation column names are, by default, identical to the logical column names in the Business Model and Mapping layer. However, you can present a different name by clearing both the **Use Logical Column Name** and the **Custom display name** options in the Presentation Column dialog.

To provide a convenient organization for your users, you can drag and drop a column from a single logical table in the Business Model and Mapping layer onto multiple presentation tables. This lets you create categories that make sense to the users. For example, you can create several presentation tables that contain different classes of measures: one containing volume measures, one containing share measures, one containing measures from a year ago, and so on.

Caution: When you drag columns to presentation tables, make sure that columns with the same name or an alias of the same name do not already exist.

Typically, presentation columns are created automatically by dragging and dropping logical columns from the logical layer.

To edit the properties of a presentation column:

1. In the Presentation layer, double-click a presentation column to display the Presentation Column dialog.
2. In the General tab, to specify a name that is different from the Logical Column name, clear **Use Logical Column Name**, and then type a name for the column. Note that aliases are created automatically whenever presentation objects are renamed, so that any queries using the original name do not break.
3. To set permissions for this presentation column, click **Permissions**. See ["Setting Permissions for Presentation Layer Objects"](#) for more information.
4. Select **Custom display name** or **Custom description** to dynamically display a custom name or custom description based on a session variable, typically for localization purposes. See ["Localizing Oracle Business Intelligence"](#) in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization.

Note: You can also use the **Custom display name** and **Custom description** fields to propagate UI hints (labels and tooltips) from an ADF data source to display in Oracle BI Answers. See ["Propagating Labels and Tooltips from ADF Business Component Data Sources"](#) for more information about how to set up ADF data source UI hints.

5. The **Logical Column** field displays the name of the logical column for this presentation column. Click **Edit** to make any changes to the logical column object.
6. Use the Aliases tab to specify or delete aliases for this presentation column. See ["Creating Aliases \(Synonyms\) for Presentation Layer Objects"](#) for more information about aliases.

To reorder a presentation column:

1. In the Presentation layer, right-click a presentation table and select **Properties**.
2. Click the **Columns** tab.
3. Select the column you want to reorder.
4. Use drag-and-drop to reposition the column, or click the **Up** and **Down** buttons.
5. Click **OK**.

Working with Presentation Hierarchies and Levels

Presentation hierarchies and presentation levels provide an explicit way to expose the multidimensional model in Oracle BI Answers. When presentation hierarchies and levels are defined in the Presentation layer, roll-up information is displayed in the Oracle BI Answers navigation pane, providing users with important contextual information.

Be aware that members in a presentation hierarchy are not visible in the Presentation layer. Instead, you can see hierarchy members in Answers.

Most importantly, users can create hierarchy-based queries using these objects. Presentation hierarchies expose analytic functionality such as member selection, custom member groups, and asymmetric queries.

As with other Presentation layer objects, you can also provide localization information and apply fine-grained access control to presentation hierarchies and levels.

If you have a repository from a previous release, note that presentation hierarchies do not appear in the Presentation layer automatically as part of the RPD upgrade process. You must manually create these objects by dragging logical dimensions from the Business Model and Mapping layer to the appropriate presentation tables.

This section contains the following topics:

- [Creating and Managing Presentation Hierarchies](#)
- [Creating and Managing Presentation Levels](#)

Creating and Managing Presentation Hierarchies

To create a presentation hierarchy, you can drag a logical dimension hierarchy from the Business Model and Mapping layer to a table in the Presentation layer. The presentation hierarchy object must be located within a presentation table, unlike in the Business Model and Mapping layer, where logical dimensions are peer objects of tables. Presentation hierarchies are also displayed within their associated tables in Oracle BI Answers, providing a conceptually simpler model.


If a logical dimension spans multiple logical tables in the Business Model and Mapping layer, it is a best practice to model the separate logical tables as a single presentation table in the Presentation layer.

There are several ways to create presentation hierarchies:


- When you drag an entire business model to the Presentation layer, the presentation hierarchies and constituent levels appear automatically, along with other presentation objects.
- When you drag a logical dimension table to the Presentation layer, presentation hierarchies and levels based on those dimensions are created automatically.
- You can also drag individual logical dimensions to the appropriate presentation tables to create corresponding presentation hierarchies within those tables.
- As with most other objects in the Administration Tool, you can right-click a presentation table and select **New Object > Presentation Hierarchy** to manually define the object.

You can also drag an individual logical level from the Business Model and Mapping layer to a presentation table to create a presentation hierarchy that is a subset of the logical dimension hierarchy.

For example, suppose a logical dimension has the levels All Markets, Total US, Region, District, Market, and Market Key. Dragging and dropping the entire logical dimension to the corresponding presentation table appears as follows:

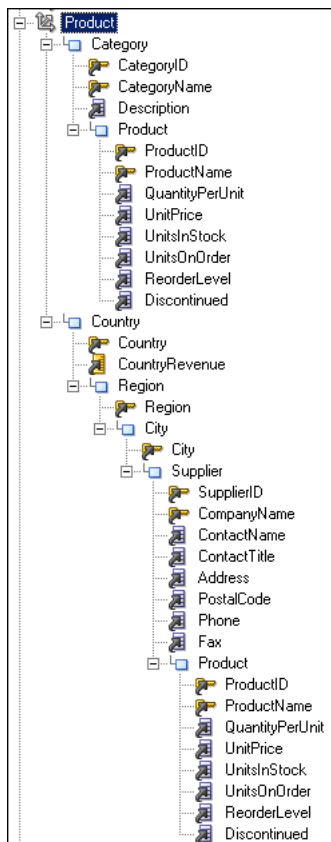


However, dragging and dropping the Region level to the same presentation table appears as follows:




Modeling Dimensions with Multiple Hierarchies in the Presentation Layer

For logical dimensions that contain multiple logical hierarchies, multiple separate presentation hierarchies are created. For example, the following logical dimension called Product contains the two hierarchies Category and Country:



In the Business Model and Mapping layer, this logical dimension is modeled as a single dimension object that contains multiple hierarchies. In contrast, the Presentation layer models this dimension as two separate objects: one that displays the drill path through the Category level, and another that shows the drill path through the Country level, as follows:



Editing Presentation Hierarchy Objects

You can edit presentation hierarchy properties, including setting permissions to apply role-based access control, setting a custom display name for localization purposes, and changing the levels in a hierarchy.

To edit the properties of a presentation hierarchy:

1. In the Presentation layer, double-click a presentation hierarchy to display the Presentation Hierarchy dialog.
2. In the General tab, you can change the following:
 - **Name.** Note that aliases are created automatically whenever presentation objects are renamed, so that any queries using the original name do not break.
 - **Permissions.** See ["Setting Permissions for Presentation Layer Objects"](#) for more information.
 - **Custom display name** and **Custom description.** Select **Custom display name** or **Custom description** to dynamically display a custom name or custom description based on a session variable, typically for localization purposes. See *"Localizing Oracle Business Intelligence"* in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization.
 - **Logical Dimension.** This field displays the name of the logical dimension for this presentation hierarchy. Click **Browse** to select a different logical dimension.
3. The Levels tab lists the levels within the hierarchy and their order. This tab is not available for parent-child hierarchies. You can add, delete, or reorder levels. You can also click the **Edit** button to edit properties for a particular level. See ["Creating and Managing Presentation Levels"](#) for information about level properties.
4. The Display Columns tab is only available for parent-child hierarchies. Because parent-child hierarchies do not contain levels, display columns are defined for the presentation hierarchy object as a whole. Use the Display Columns tab to define which columns should be used for display for this parent-child hierarchy.

You can add, delete, or reorder display columns. You can also click the **Edit** button to edit properties for a particular column.
5. Use the Aliases tab to specify or delete aliases for this presentation hierarchy. See ["Creating Aliases \(Synonyms\) for Presentation Layer Objects"](#) for more information about aliases.

Creating and Managing Presentation Levels

Presentation levels are displayed within hierarchical columns in Oracle BI Answers. Presentation levels are typically created automatically when presentation hierarchies are created.

To edit the properties of a presentation level:

1. In the Presentation layer, double-click a presentation level to display the Presentation level dialog.
2. In the General tab, you can change the following:
 - **Name.** Note that aliases are created automatically whenever presentation objects are renamed, so that any queries using the original name do not break.
 - **Permissions.** See ["Setting Permissions for Presentation Layer Objects"](#) for more information.
 - **Custom display name and Custom description.** Select **Custom display name** or **Custom description** to dynamically display a custom name or custom description based on a session variable, typically for localization purposes. See *"Localizing Oracle Business Intelligence"* in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about localization.
 - **Logical Level.** This field displays the name of the logical level for this presentation level. Click **Browse** to select a different logical level.
3. The Drill To Levels and Drill From Levels tabs are reserved for a future release and are not currently used. An additional option called **Generate Drill Graph**, available as a right-click option for any Presentation layer object, is also reserved for a future release.
4. Use the Display Columns tab to define which columns should be used for display for that level (on drill-down). For example, if two columns called "Name" and "ID" exist at the same level, you can choose to display "Name" because it is the more user-friendly option.

You can add, delete, or reorder display columns. You can also click the **Edit** button to edit properties for a particular column.

As an alternative to defining display columns in this tab, you can drag a presentation column directly onto the presentation level in the Presentation layer of the Administration Tool. Doing this automatically adds the column as a display column for the presentation level.

Note that the display columns that appear by default when a presentation level is created are based on which key columns for the corresponding logical level have the **Use for display** option selected.

5. Use the Aliases tab to specify or delete aliases for a presentation level. See ["Creating Aliases \(Synonyms\) for Presentation Layer Objects"](#) for more information about aliases.

Setting Permissions for Presentation Layer Objects

You can apply access control to restrict which individual users or application roles (groups) can access particular presentation layer objects. For example, you can provide read-only access to a set of presentation tables for a particular application role, read-write access for a second application role, and no access for a third application role.

You can also use the Identity Manager to set up privileges and permissions. The Identity Manager is useful for setting permissions for individual application roles to many objects at once, unlike permissions in the Presentation layer, which you can only set for one object at a time. See ["Setting Up Object Permissions"](#) for information about

setting up object permissions in the Identity Manager. For a full description of data access security in Oracle Business Intelligence, see [Chapter 13](#).

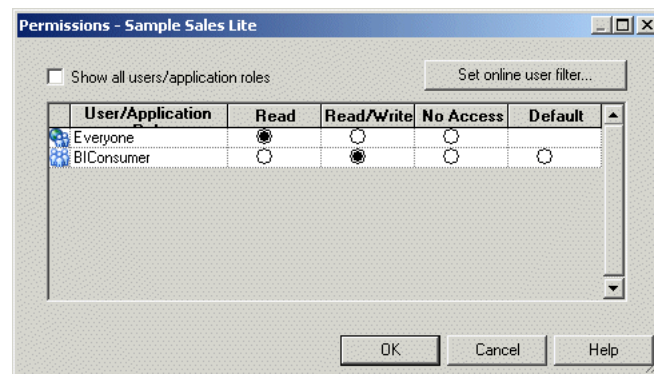
You can control what level of privilege is granted by default to users and application roles for repository objects without explicit permissions set. To do this, set the `DEFAULT_PRIVILEGES` parameter in the `NQSCONFIG.INI` file. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

To set permissions for presentation layer objects:

1. In the Presentation layer, double-click a presentation object, such as a subject area, table, column, or hierarchy.
2. In the General tab, click **Permissions**.

Figure 11–2 shows the Permissions dialog.

Figure 11–2 Permissions Dialog



3. In the Permissions dialog, any users or application roles with the **Default** permission do not appear in the User/Application Roles list. Select **Show all users/application roles** to see users and application roles with the Default permission.

In online mode only, by default, no users are retrieved, even when **Show all users/application roles** is selected. Click **Set online user filter** to specify the set of users you want to retrieve.

The filter is empty by default, which means that no users are retrieved. Enter `*` to retrieve all users, or enter a combination of characters for a specific set of users, such as `A*` to retrieve all users whose names begin with the letter A. The filter is not case-sensitive.

4. For each user and application role, you can allow or disallow access privileges for this presentation object by selecting one of the following options:
 - **Read.** Only allows read access to this object.
 - **Read/Write.** Provides both read and write access to this object.
 - **No Access.** Explicitly denies all access to this object.
 - **Default.** No object-specific access rules apply. The default access control for this user or application role, as defined in the Identity Manager, controls what the user or application role can do with this object.

Note that any permissions applied to a presentation object also apply to its child objects. For example, permissions applied to a presentation table also apply to any presentation columns, hierarchies, and levels in that table.

5. Click **OK**.
6. Click **OK** in the Properties dialog for this presentation object.

Generating a Permission Report for Presentation Layer Objects

You can generate a permission report for individual presentation layer objects to see a summary of how permissions have been applied for that object. To do this, right-click any presentation object and select Permission Report. The Permission Report dialog displays the name and a description of the presentation object, along with a list of users/application roles and their permissions.

Sorting Columns in the Permissions Dialog

There are six ways that you can sort the types and User/Application Role names in the Permissions dialog. To change the sort, click the heading of the first or second column. The first column has no heading and contains an icon that represents the type of user or application role. The second column contains the name of the User/Application Role object. Note that you cannot sort on the columns for individual object permissions (like Read, Read/Write, and so on).

There are three ways to sort by type, and two ways to sort the list of user and application role names. This results in a total of six possible sort results ($3 \times 2 = 6$). The following list shows the sort results available by clicking the type column:

- Everyone, Application Roles, Users (ascending by name of type)
- Users, Application Roles, Everyone (descending by name of type)
- Type column is in no particular order (Type value is ignored, as all names in User/Application Role column are sorted in ascending order by value in User/Application Role column)

The following list shows the sort results available by clicking the User/Application Role column:

- Ascending within the type
- Descending within the type

Creating Aliases (Synonyms) for Presentation Layer Objects

Each presentation object can have a list of aliases (synonyms) for its name that can be used in Logical SQL queries. To create the list of aliases, use the Alias tab in the Properties dialog for the appropriate presentation object (subject area, presentation table, presentation hierarchy, presentation level, or presentation column).

Because Presentation layer objects are often deleted and then re-created during the repository development process, it is best to wait until your logical business model is relatively stable before creating aliases for presentation objects.

You can use this feature to rename presentation objects without breaking references that any existing requests have to the old names, including requests from Answers, Oracle BI Publisher, or other Logical SQL clients. If you are still developing a new repository, you might want to wait until the repository is stable before renaming objects.

For example, consider a subject area called "Sample Sales Reduced" that contains a presentation table called "Facts Other." If you rename the presentation column called "# of Customers" to "Number of Customers," any requests that use "# of Customers" fail. However, if you add "# of Customers" to the list of synonyms in the Alias tab for the "Number of Customers" column, then queries containing both "# of Customers" and "Number of Customers" succeed and return the same results.

Note the following:

- Aliases for presentation objects do not appear in Answers or other query clients when creating new queries. Only the primary names of subject areas, hierarchies, levels, tables, and columns appear.
- This feature works in a different way from SQL aliases or the alias feature in the Physical layer. It simply provides synonyms for object names, much like "synonyms" in SQL.
- Aliases are created automatically when you rename presentation objects. For example, if you change Catalog to Catalog1, the original name Catalog is added to the Aliases list.
- You cannot rename a Presentation layer object to a name that is already in use as an alias for an object of the same type.

To add or delete an alias for a presentation object:

1. In the Presentation layer, double-click a presentation object, such as a subject area, table, column, or hierarchy.
2. Click the Aliases tab.
3. To add an alias, click the **New** button, and then type the text string to use for the alias.
4. To delete an alias, select the alias you want to delete from the **Aliases** list, then click the **Delete** button.
5. Click **OK**.

Creating and Persisting Aggregates for Oracle BI Server Queries

Most data warehouse practitioners create aggregated data tables to dramatically improve the performance of highly summarized queries. These aggregate tables store precomputed results that are aggregated measures (typically summed) over a set of dimensional attributes. Using aggregate tables is a typical technique used to improve query response times in decision support systems.

If you write SQL queries or use a tool that only understands what physical tables exist and not their meaning, then using aggregate tables becomes more complex as the number of aggregate tables increases. The aggregate navigation capability of the Oracle BI Server allows queries to use the information stored in aggregate tables automatically. The Oracle BI Server lets you concentrate on asking the right business question, and then the server decides which tables provide the fastest answers.

Oracle Business Intelligence has an aggregate navigation feature to take advantage of those aggregates in source databases (for more information, see [Chapter 10](#)). However, it can be time consuming to create and maintain the data aggregation, as well as load database scripts and the corresponding metadata mappings. For that reason, Oracle Business Intelligence provides an aggregate persistence feature that automates the creation and loading of the aggregate tables and their corresponding Oracle Business Intelligence metadata mappings.

This chapter explains how to set up and use aggregate persistence in Oracle Business Intelligence.

This chapter contains the following topics:

- [About Aggregate Persistence in Oracle Business Intelligence](#)
- [Identifying Query Candidates for Aggregation](#)
- [Using the Aggregate Persistence Wizard to Generate the Aggregate Specification](#)
- [Writing the Create Aggregates Specification Manually](#)
- [Running the Aggregate Specification Against the Oracle BI Server](#)
- [Troubleshooting Aggregate Persistence](#)

About Aggregate Persistence in Oracle Business Intelligence

Use the Aggregate Persistence feature to create aggregates for Oracle BI Server queries. The Aggregate Persistence Wizard lets you automate the creation of the aggregate specification script. When you run this script against a live Oracle BI Server, aggregate tables are created and are mapped into the metadata for navigation. When

aggregates are persisted, indexes and statistics are created on relational tables for greater performance.

The Aggregate Persistence Wizard creates a SQL script that you can run on a scheduled basis against the Oracle BI Server. In the Aggregate Persistence Wizard, you specify the measures, dimensionality, and other parameters of each star or cube based on your performance design. The script should run after each load of the base-level tables, so that the aggregates are always synchronized with the detail-level data when the load window completes and users begin to run queries.

Aggregate creation runs against the master server in a cluster. It takes some time for the metadata changes to propagate to the slaves. The cluster refresh time is a user-controlled option and results might be incorrect if a query hits a slave server before it is refreshed. It is the administrator's responsibility to set an appropriate cluster refresh interval.

Aggregate persistence requires a dedicated connection pool to create tables or cubes in the target database that will hold the aggregates. Because the Oracle BI repository enables federation, the aggregated target can be on the same database as the detailed source, or in a completely different database. This dedicated connection pool must be created before you run the Aggregate Persistence Wizard, so it can be selected during the appropriate step of the wizard.

The default prefix `SA_` is automatically added to dimension (level) aggregates. You can change this default prefix by updating the `AGGREGATE_PREFIX` parameter in the `AGGREGATE_PERSISTENCE` section of the `NQSCONFIG.INI` file:

```
AGGREGATE_PREFIX = "prefix_name" ;
```

The target schema used to store aggregates must be appropriately secured and should not allow public access. The schema should have privileges to connect, create, and drop tables and indexes. By default, only users who belong to the `BIAdministrators` group can manage aggregates.

Do not use aggregate persistence against tables with active Virtual Private Database (VPD) security filters. There is a possibility that the aggregate information might be persisted without the VPD filter, posing a security risk.

Identifying Query Candidates for Aggregation

When creating aggregates, you must identify which queries would benefit substantially from aggregated data. You will achieve the best results by aggregating to the highest level possible. To identify slow-running queries, perform the following tasks:

- **Enable usage tracking in the Oracle BI Server.** Usage tracking statistics can be used in a variety of ways, such as database optimization, aggregation strategies, and billing users or departments based on the resources they consume. The Oracle BI Server tracks usage at the detailed query level. When you enable usage tracking, statistics for every query are written to a usage tracking log file or inserted into a database table.

Note: It is strongly recommended that you use the direct insertion into a database method for usage tracking. See "Managing Usage Tracking" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for full information about usage tracking.

- **Analyze the query run times and identify the slowest running queries as candidates for aggregation.** The run time for creating aggregates is dependent on the type of aggregates selected by the user. Creating aggregates from large fact tables is slower than from smaller tables. You should carefully select the aggregates to be created.

Using the Aggregate Persistence Wizard to Generate the Aggregate Specification

You can use the Aggregate Persistence Wizard to create the SQL file that will be used to create and load aggregate tables and map them into the metadata. The resulting SQL file must be executed against a running Oracle BI Server.

Note: It is strongly recommended that you use the Aggregate Persistence Wizard because it automatically enforces many of the constraints necessary when generating the aggregate specification. However, you can manually write the aggregate Logical SQL as an alternative to using the wizard. Make sure to follow the guidelines described in "[Writing the Create Aggregates Specification Manually](#)" if you choose to write your own aggregates specification.

To use the Aggregate Persistence Wizard:

1. In the Administration Tool, select **Tools > Utilities > Aggregate Persistence**, and then click **Execute**.
2. On the Select File Location screen, specify the complete path and file name of the aggregate creation script. You can specify a new or an existing file name.

Typically, when you run the SQL script against the Oracle BI Server, it creates DDL and runs it against the target database schema to create the aggregate tables, then loads them from the source, and finally creates the Oracle BI Server metadata so the aggregate navigation feature can use the new tables.

Alternatively, you can select **Generate target DDL in a separate file** if you want the DDL to be stored in a separate file from the Oracle BI Server SQL script. When you select this option, two SQL scripts are generated:

- The create aggregates script (*script_name*)
- The prepare aggregates script (*script_name_DDL*)

Both files are stored in the following location:

```
ORACLE_INSTANCE\bifoundation\OracleBIServerComponent\coreapplication_obisn\
aggr
```

Selecting **Generate target DDL in a separate file** gives you the flexibility to alter the auto-generated DDL and run it independently of the Oracle BI Server. For example, you may want to alter the storage parameter or index settings. When you select this option, you first make manual updates to the DDL file, then you run the DDL file (prepare aggregates), then you run the create aggregates script.

Click **Next** after you have finished specifying options on the Select File Location screen.


3. In the Select Business Measures screen, select the measures on which you want to aggregate. To do this, select a business model in the upper pane, then select a

single fact table or a set of measures in the lower pane. You cannot select measures that span multiple fact tables. Use Ctrl-click to select multiple measures, or use Shift-click to select a range of consecutive measures.

Note that the **View Script** button is not available during the creation of the first aggregate table block.

Figure 12–1 shows the Select Business Measures screen.

Figure 12–1 Aggregate Persistence Wizard: Select Business Measures Screen



Click **Next** after you have selected the appropriate measures.

4. In the Select Levels screen, specify the level of aggregation by selecting a logical level for one or more dimensions. You can specify a surrogate key to be used for the fact-dimension join.

The default join option between the aggregated fact and dimension tables is the primary key defined in the logical level you selected. If the primary key of the level is large and complex, the join to the fact table is expensive, so using a surrogate key is recommended in this case. A surrogate key is an artificially generated key, usually a number. For example, a surrogate key in the level aggregate table would simplify this join, removing unnecessary (level primary key) columns from the fact table and resulting in a leaner fact table.

Using a surrogate key only changes the query response time, not the logical results of the queries. However, generating the surrogate keys can have the side effect of increasing the aggregate table load time. Therefore, the recommended setting is as follows:


- If the primary key for the logical level you have selected is already a single, numeric column, you typically should not select the **Use Surrogate Key** option since it may add to load processing time without producing a performance benefit.
- If the primary key for the logical level you have selected is a text string, or consists of multiple logical columns, you typically should use a surrogate key

to improve the performance of the queries that join to that aggregate dimension. However, keep in mind that generating the surrogate key can increase the load time for that aggregate dimension table.

See ["Adding Surrogate Keys to Dimension Aggregate Tables"](#) for additional information about surrogate keys.

Figure 12–2 shows the Select Levels screen.

Figure 12–2 Aggregate Persistence Wizard: Select Levels Screen




Click **Next** after you have selected the appropriate level of aggregation.

- In the Select Connection Pool screen, select the appropriate items to specify a location for the aggregate table.

A default aggregate table name is provided, and a prefix is added to the table name. The default prefix for the generated fact table is `ag`. For tables created for dimension (level) aggregates, the default prefix is `SA_` and can be changed by updating the `AGGREGATE_PREFIX` property in `NQSConfig.INI`. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about changing configuration settings.

Figure 12–3 shows the Select Connection Pool screen.

Figure 12-3 Aggregate Persistence Wizard: Select Connection Pool Screen

Click **Next** after you have provided connection pool information.

6. In the Finish screen, the **View Script** button becomes available for use, and the Logical SQL script appears for your review. Choose whether to define another aggregate (default) or end the wizard, and then click **Next**.
7. In the Finish Script screen, the complete path and file name appears. Click **Finish**.

See "[Running the Aggregate Specification Against the Oracle BI Server](#)" for information about using the SQL file to create aggregate tables.

Writing the Create Aggregates Specification Manually

If you choose not to use the Aggregate Persistence Wizard to create the script file, you can write the file manually. It is recommended that you use the Aggregate Persistence Wizard.

If you do not want the Oracle BI Server to modify your databases during aggregate creation, then you can specify this in the Aggregate Persistence Wizard by selecting the option **Generate target DDL in a separate file**. The Aggregate Persistence Wizard will create a DDL file (the "prepare aggregates" script) that you can use to create the empty aggregate tables. After this, you need to run the "create aggregates" script to populate the aggregate tables. This option provides some flexibility in case the database access to create tables is restricted. Note that you must run the prepare aggregates script before you run the create aggregates script.

This section contains the following topics:

- [What Constraints Are Imposed During the Create Process?](#)
- [How to Write the Create Aggregates Specification](#)
- [Adding Surrogate Keys to Dimension Aggregate Tables](#)

What Constraints Are Imposed During the Create Process?

The following constraints are imposed during the create process:

- **Valid measures.** A valid measure must have a valid aggregation rule. The following constraints apply to level-based measures:
 - If the level is grand total alias, then that dimension must not be present in the list of levels for that aggregate specification.
 - Any other level defined for this measure must be present in the list of levels for that aggregate specification.

If the above constraints are not met, then the entire aggregate specification is discarded. In addition, a measure is ignored by the create process if any of the following conditions are true:

- Measure is mapped to a session or repository variable.
- Measure is a derived measure.
- Measure has a default aggregation rule of None.

Measures that are ignored do not necessarily affect the aggregate specification. The remaining measures are used to create the aggregate.

- **Valid levels.** A valid level must have a valid primary key. If a level is invalid, the aggregate specification is discarded. Also, attributes of a level or its primary key are ignored if any of the following conditions are true:
 - Attribute is mapped to session or repository variables.
 - Attributes are not from the same logical table.
- **Valid aggregate specification.** A valid aggregate specification has the following properties:
 - Name length is between 1 and 18 characters (inclusive).
 - At least one valid level must be specified.
 - At least one valid measure must be specified.
 - Must have a valid connection pool.
 - Must have a valid output container (database/catalog/schema).
 - Connection pool and container must belong to the same database.
 - Only one level per dimension can be specified.
 - Measures can only be from the same fact table.
 - All logical components of the specification must be from the same subject area.

An aggregate specification is ignored if the name already exists in the output container because level aggregates are scoped by the entire database. However, if different catalogs or schemas are specified for the same fact aggregate name, it is allowed to have multiple facts with the same name but different scope in the same database.

Note that the aggregate specification is discarded if any dimension is not joined to a fact.

How to Write the Create Aggregates Specification

All metadata names (except for logical fact columns) are fully qualified. There are two modes of operation: Create and Delete. It is strongly recommended that you place all aggregate specifications under a single Create Aggregates statement.

Follow these guidelines when writing the aggregate specification:

- Begin the script file with a Delete statement. It is essential to delete system-generated aggregates before creating new ones. This ensures that data is consistent and removes invalid or incomplete aggregates before you run the Create operation. The following statement is the syntax for deleting aggregates:

```
Delete aggregates;
```

- The next statement should be a Create statement. The following is the syntax for creating aggregates:

```
Create|Prepare aggregates
aggr_name_1
for logical_fact_table_1 [(logical_fact_column_1, logical_fact_column_2,...)]
at levels (level_1, level_2, ...)
using connection pool connection_pool_name_1
in schema_name_1
[ ,aggr_name_2
for logical_fact_table_3 [(logical_fact_column_5, logical_fact_column_2,...)]
at levels (level_3, level_2, ...)
using connection pool connection_pool_name_2
in schema_name_2] ;
```

- To specify multiple aggregates in a single Create Aggregates statement, follow these guidelines:
 - Ensure that each of the multiple aggregate specifications are separated by a comma, and the entire aggregate creation script is terminated with a semicolon.
 - In this file, only one Delete Aggregates statement should be specified at the beginning. Make sure that only one delete is issued per ETL run (unless a reset is called for).

Caution: Any aggregate scripts that are run after the first one should not have a Delete Aggregates statement, or all previously created aggregates are removed.

For information about creating aggregates with surrogate keys, see the following section.

Adding Surrogate Keys to Dimension Aggregate Tables

The join option default between fact and level aggregate tables uses primary keys from the level aggregate. If the primary key of the level is large and complex (composite of many columns), then the join to the fact table is expensive. A surrogate key is an artificially generated key, usually a number. A surrogate key, in the level aggregate table, simplifies this join and removes unnecessary columns (level primary key) from the fact table, resulting in a smaller-sized fact table. Adding surrogate keys to the dimension (level) aggregate tables can simplify joins to the fact tables and might improve query performance. Additionally, a surrogate key makes sure that each aggregate table has a unique identifier.

There might be cases in which a level is shared among multiple fact tables. One fact might use surrogate keys, and another might use primary keys from the dimension aggregate. The following are some options for resolving this issue:

- Set a metadata property for levels that indicates whether to use surrogate keys or primary keys.
- Always create a surrogate key for a level aggregate (relatively low cost operation). Then, decide later if the fact aggregate should join to it using a surrogate or primary key.

An alternative to specifying the join type for each dimension is to specify if surrogate keys should be used for the entire star. This would result in simpler syntax, but would also restrict the available user options and slow the aggregate creation process.

About the Create/Prepare Aggregates Syntax

The following syntax for create/prepare aggregates contains the change for [Using_Surrogate_Key]. The surrogate key option can be specified for each level. If unspecified, the fact and dimension tables are joined using the primary key from the level aggregate.

```
Create|Prepare aggregates
aggr_name_1
[file output_file_name]
for logical_fact_table_1 [(logical_fact_column_1, logical_fact_column_2,...)]
at levels (level_1 [Using_Surrogate_Key], level_2, ...)
using connection pool connection_pool_name_1
in schema_name_1
[ ,aggr_name_2
for logical_fact_table_3 [(logical_fact_column_5, logical_fact_column_2,...)]
at levels (level_3, level_2, ...)
using connection pool connection_pool_name_2
in schema_name_2] ;
```

About Surrogate Key Output from Create/Prepare Aggregates

The changes to the current process are restricted to the physical metadata layer in the repository and the database.

When you use the Using_Surrogate_Key join option, the following describes the results:

- For a level aggregate, the following occurs:
 - In the physical metadata, the following occurs:
 - * The level aggregate table has a new column called *levelName_upgradeIDSK* (check for collisions). This is the surrogate key column for the dimension aggregate. Note that *levelName* is truncated if the total number of characters exceeds 18.
 - * The type of this column is SMALLUINT.
 - In the database, the following occurs:
 - * The level aggregate table also has a corresponding column called *levelName_upgradeIDSK*. Again, *levelName* is truncated if the total number of characters exceeds 18.
 - * It can be populated using RCOUNT().
- For a fact aggregate, the following occurs:

- In the physical metadata, the following occurs:
 - * The fact aggregate table no longer contains columns from the level's primary keys.
 - * Instead, a new column that corresponds to the level aggregate's surrogate key is added to the table.
 - * The type of this column is identical to the level's surrogate key.
 - * The column has the same name as that in the level aggregate (check for collisions).
 - * The fact table and the level table are joined using this surrogate key only.
- In the database, the following occurs:
 - * The fact aggregate table also has the corresponding surrogate key.
 - * It is populated using new capabilities to be available through Populate.

Running the Aggregate Specification Against the Oracle BI Server

Before you run the aggregate specification against the Oracle BI Server, you should set an appropriate logging level. Trace logs are logged to `nquery.log` if the logging level is at least 2. The logging events include the aggregate execution plan and the order in which the aggregates are created and deleted. Higher logging levels provide more details about the query and execution plans. Error logs are logged to `nquery.log` if the logging level is at least 1, and to `nqserver.log` regardless of the logging level.

After generating the SQL script file and setting the logging levels, you can run the SQL script using `nqcmd`. The SQL file must be executed against a running Oracle BI Server. See "[Testing and Refining the Repository](#)" for more information about running `nqcmd`.

After executing the SQL script, aggregates are created and persisted in the Oracle BI Server metadata, as well as in the back-end databases.

Troubleshooting Aggregate Persistence

The following is a list of some reasons errors can occur:

- Network failure.
- No disk space on the database.
- Bad aggregate request.

If there is an error in the creation of any aggregate, then the entire aggregate request is terminated and subsequent aggregates are not created. Aggregates that are already created and checked in remain checked in. If there are errors, you must remove them at the time of the error or at the next ETL run in one of the following ways:

- Manually remove the aggregates from the metadata and the database.
- Automatically remove all the aggregates using the Delete Aggregates specification.

Applying Data Access Security to Repository Objects

Data access security controls rights to view and modify data. You can use several different methods of data access security with Oracle Business Intelligence: row-level security (implemented either in the repository or in the database), object permissions, and query limits. This chapter provides information about the different types of data access security and explains how to set them up.

Other security tasks, including setting up SSL connections, managing users, groups, and application roles, setting up custom LDAP servers, and managing custom authenticators, are covered in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*. Note that you must create users and application roles before you can implement data access security.

You should plan to implement data access security in the Administration Tool in online mode. If you must perform data access security tasks in offline mode, be sure to read "[About Applying Data Access Security in Offline Mode](#)" first.

Data access security auditing is covered by the Oracle Business Intelligence usage tracking feature. See "Managing Usage Tracking" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

This chapter contains the following topics:

- [About Data Access Security](#)
- [Setting Up Row-Level Security](#)
- [Setting Up Object Permissions](#)
- [Setting Query Limits](#)
- [About Applying Data Access Security in Offline Mode](#)
- [About the List of Users in the Administration Tool](#)

About Data Access Security

After developing your metadata repository, you need to set up your data security architecture to control access to source data. Data access security accomplishes the following goals:

- To protect business data queried from databases
- To protect your repository metadata (such as measure definitions)
- To prevent individual users from hurting overall system performance

Oracle Business Intelligence supports three types of data security: row-level security, object permissions, and query limits (governors). Object permissions and query limits are set up in the repository and are enforced only by the Oracle BI Server. Row-level data security, however, can be implemented and enforced in both the repository, and in the database.

Even if you choose to implement row-level security in the database, you should still set up object permissions and query limits in the repository. Although it is possible to provide database-level object restrictions on individual tables or columns, objects to which users do not have access are still visible in all clients, even though queries against them will fail. It is better to set up object permissions in the repository, so that objects to which users do not have access are hidden in all clients.

Because a variety of clients can connect to the Oracle BI Server, you cannot implement or enforce data security in Oracle BI Presentation Services. Oracle BI Presentation Services provides an extensive set of security controls that let you set up privileges to access functionality in the Oracle Business Intelligence user interface, as well as dashboards and analyses objects. However, Oracle BI Presentation Services does not provide data access security. If you only implement security controls in Oracle BI Presentation Services, you will be exposed to SQL injection hacker attacks and other security vulnerabilities. You must provide object-level security in the repository to create rules that apply to all incoming clients.

See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for more information about the security controls available in Oracle BI Presentation Services.

Where Do I Find Information About Security Tasks?

Oracle Business Intelligence security tasks are covered in this guide, in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*, and in other sources. [Table 13–1](#) summarizes the Oracle Business Intelligence security tasks and where to go for more information.

Table 13–1 Security Tasks in Oracle Business Intelligence

Task	Location
Setting up user authentication with the default authentication provider or an alternative authentication provider	"Managing Security Using the Default Security Configuration" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Creating and managing users and groups in the default authentication provider	"Creating and Managing Users and Groups in the Embedded WebLogic LDAP Server" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Creating application roles and managing policies in the default policy store	"Managing the Policy Store" in <i>Oracle Fusion Middleware Application Security Guide</i>
Viewing and understanding the default Oracle Business Intelligence permissions used with application roles in the policy store	"Default Permissions" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing system credentials in the default credential store	"Default Credentials" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>

Table 13–1 (Cont.) Security Tasks in Oracle Business Intelligence

Task	Location
Applying data access security in offline mode and setting up placeholder application roles	"About Applying Data Access Security in Offline Mode"
Setting up row-level data security	"Setting Up Row-Level Security"
Setting repository object permissions	"Setting Up Object Permissions"
Setting query limits (governors)	"Setting Query Limits"
Viewing users in the Administration Tool	"About the List of Users in the Administration Tool"
Setting up single sign-on (SSO)	"Enabling SSO Authentication" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Enabling SSL communication	"SSL Configuration in Oracle Business Intelligence" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>
Managing custom authenticators	"Authenticating by Using a Custom Authenticator Plug-In" in <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>

Setting Up Row-Level Security

You can choose to set up row-level security in the repository, or in the database. Implementing row-level security in the repository provides many benefits, including the following:

- All users share the same database connection pool for better performance
- All users share cache for better performance
- You can define and maintain security rules that apply across many federated data sources

Implementing row-level security in the database, in contrast, is good for situations where multiple applications share the same database. Note that even when you design and implement row-level security in the database, you should still define and apply object permissions in the repository.

Although it is possible to set up row-level security in both the repository and in the database, you typically do not enforce row-level security in both places unless you have a particular need to do so.

This section contains the following topics:

- [Setting Up Row-Level Security \(Data Filters\) in the Repository](#)
- [Setting Up Row-Level Security in the Database](#)

Setting Up Row-Level Security (Data Filters) in the Repository

Data filters are a security feature that provide a way to enforce row-level security rules in the repository. Data filters are set up in the repository using the Administration Tool and are applied for a particular application role. You typically do not set up data filters if you have implemented row-level security in the database, because in this case, your

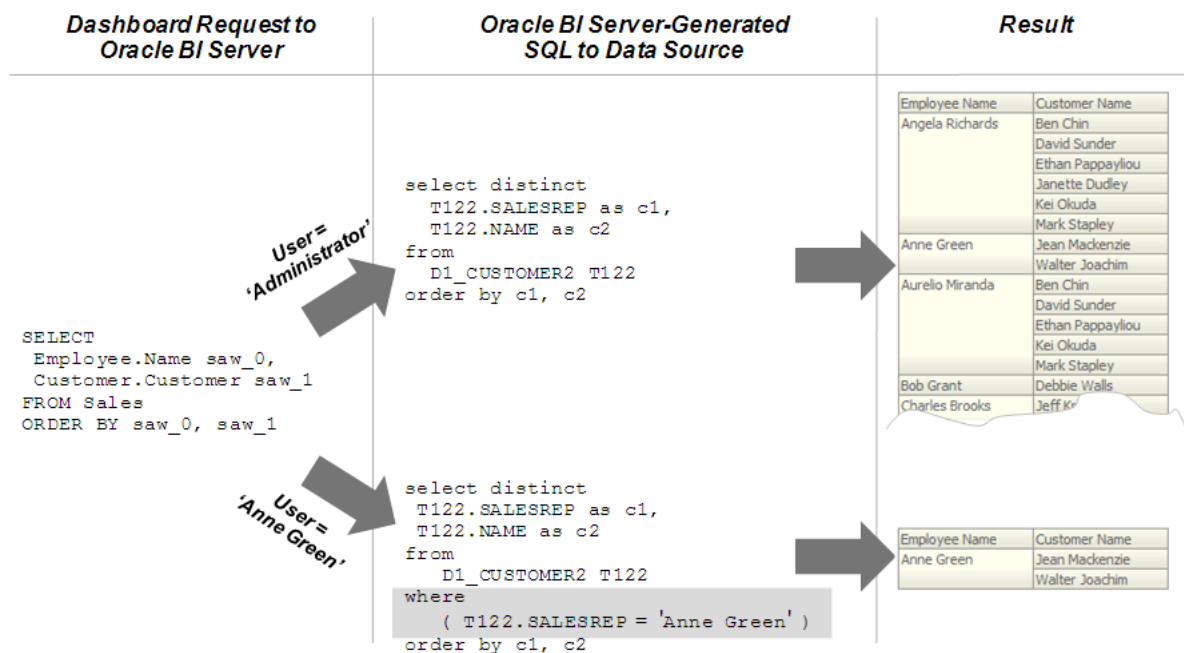
row-level security policies are being enforced by the database rather than the Oracle BI Server.

Data filters can be set for objects in both the Business Model and Mapping layer and the Presentation layer. Applying a filter on a logical object impacts all Presentation layer objects that use the object. If you set a filter on a Presentation layer object, it is applied in addition to any filters that might be set on the underlying logical objects.

Figure 13–1 illustrates how data filter rules are enforced in the Oracle BI Server. The security rules are applied to all incoming clients and cannot be breached, even when the Logical SQL query is modified.

In this example, a filter has been applied to an application role. When Anne Green, who is a member of that role, sends a request, the return results are limited based on the filter. Because no filters have been applied to the application roles for the Administrator user, all results are returned. The Oracle BI Server-generated SQL takes into account any data filters that have been defined.

Figure 13–1 Row-Level Security Enforcement in the Oracle BI Server



You should always set up data filters for particular application roles rather than for individual users.

To set up data filters to apply row-level authorization rules for queries:

1. Open your repository in the Administration Tool.
2. Select **Manage**, then select **Identity**.
3. In the Identity Manager dialog, in the tree pane, select **BI Repository**.
4. In the right pane, select the Application Roles tab, then double-click the application role for which you want to set data filters.

Note that if you are in offline mode, no application roles appear in the list unless you have first modified them in online mode. See "[About Applying Data Access Security in Offline Mode](#)" for more information.

5. In the Application Role dialog, click **Permissions**.

6. In the User/Application Role Permissions dialog, click the Data Filters tab.
To create filters, you first add objects on which you want to apply the filters. Then, you provide the filter expression information for the individual objects.
7. To add objects on which you want to apply filters, perform one of the following steps:
 - Click the **Add** button. Then, browse to locate the object you want, select it, and then click **Select**.
 - Click the **Name** field for an empty row. Then, browse to locate the object you want, select it, and then click **Select**.
8. To enter the filter expression for individual objects, perform one of the following steps:
 - Select the data filter, then click the **Expression Builder** button. Create the filter expression in Expression Builder, then click **OK**.
 - Click the **Data Filter** field for the appropriate filter, then type the filter expression.

For example, you might want to define a filter like "Sample Sales". "D2 Market". "M00 Mkt Key" > 5 to restrict results based on a range of values for another column in the table.

You can also use repository and session variables in filter definitions. Use Expression Builder to include these variables to ensure the correct syntax.
9. Optionally, select a status for each filter from the Status list. You can choose one of the following options:
 - **Enabled:** The filter is applied to any query that accesses the object.
 - **Disabled:** The filter is not used and no other filters applied to the object at higher levels of precedence (for example, through an application role) are used.
 - **Ignored:** The filter is not in use, but any other filters applied to the object (for example, through a different application role) are used. If no other filters are enabled, no filtering occurs.
10. In addition to defining new filters, you can perform other operations in the Data Filters tab. [Table 13–2](#) lists and describes the other buttons and options.

Table 13–2 Data Filters Tab: Buttons and Options






Option Name	Description
Subject Area	Select a subject area to only view data filters for that individual subject area, or select All to view all filters.
Total Filters	Lists the total number of data filters that have been defined for this particular user or application role.
Add 	Click Add to open the Browse dialog to add objects on which you want to apply data filters.
Delete 	Select a row and click Delete to remove a filter.

Table 13–2 (Cont.) Data Filters Tab: Buttons and Options

Option Name	Description
Browse 	Select a row and click Browse to change the object on which the filter is applied.
Edit Expression (Expression Builder) 	Select a row and click Edit Expression to add or change a filter expression for a particular object. You must first add an object before you can apply a filter expression to the row.
Find 	Enter text in the Find field and click Find Down or Find Up to find a particular string.

- Click **OK**, then click **OK** again to return to the Identity Manager.

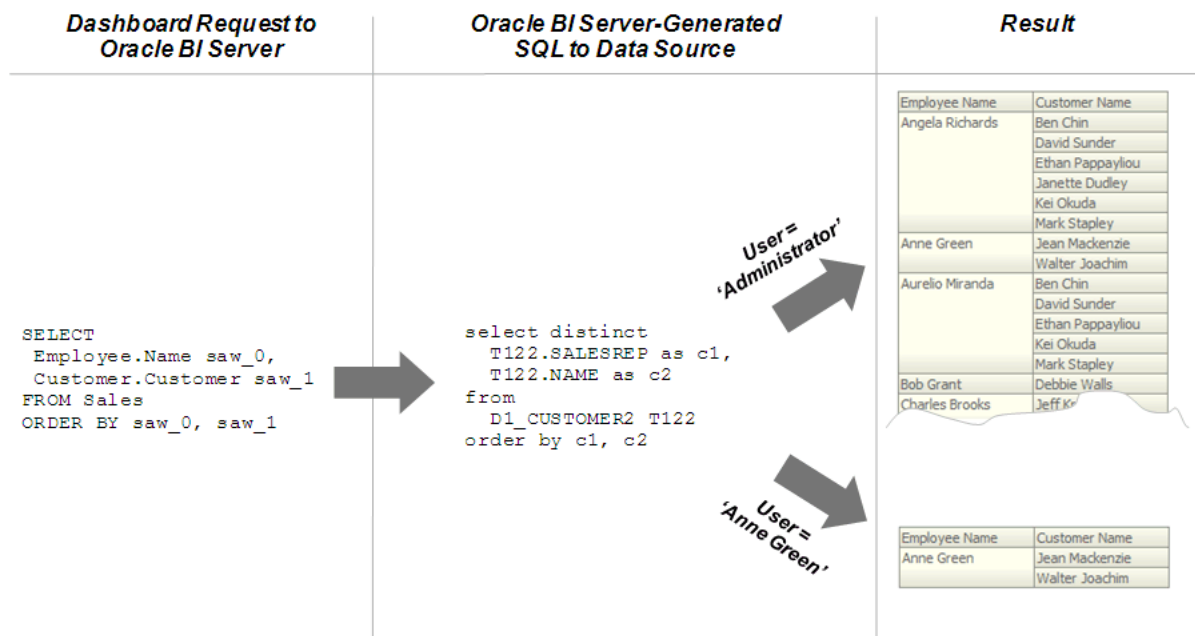
Setting Up Row-Level Security in the Database

To set up Oracle Business Intelligence for row-level security that has been implemented in the database, you can configure your connection pools so that the Oracle BI Server passes the credentials for each user to the database. The database then uses the credentials to apply its own row-level security rules to user queries.

Note that the row-level database security described in this section is different from database *authentication*, a topic discussed in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*. Rather, row-level database security provides database *authorization*. In other words, it applies access security to particular rows in the database.

[Figure 13–2](#) illustrates how row-level security is enforced in the database for Oracle Business Intelligence queries. The security rules are applied to all incoming clients and cannot be breached, even when the Logical SQL query is modified. In this example, the results returned are different depending on which user generated the query, even though the SQL query generated by the Oracle BI Server is the same. The returned results are based on rules created and enforced in the database.

Figure 13–2 Row-Level Security Enforcement in the Database




In addition to setting up Oracle Business Intelligence for row-level security in the database, you must define your set of users, permissions, and security policies in the database itself. Refer to your database documentation for more information.

To set up Oracle Business Intelligence for row-level access security in the database:

1. Open your repository in the Administration Tool.
2. Double-click the connection pool associated with the database for which you want to set up database-level security.
3. In the General tab of the Connection Pool dialog, select **Shared logon**, and then enter `:USER` and `:PASSWORD` in the **User name** and **Password** fields. The `:USER` and `:PASSWORD` syntax automatically passes the value of user credentials upon login to the database. Note that the `:USER` and `:PASSWORD` syntax does not refer to session variables.

Figure 13–3 shows the General tab of the Connection Pool dialog.

Figure 13–3 Entering Credentials for Database-Level Security in the Connection Pool



Note: Alternatively, you can use the database session context to pass end user identity to the database. Use a connection pool script to set up session context. Note that this approach does not rely on database authentication.

4. Click **OK** in the Connection Pool dialog.
5. Double-click the database object for which you want to set up database-level security.
6. In the Database dialog, select **Virtual Private Database**. Selecting this option ensures that the Oracle BI Server protects cache entries for each user.
7. Click **OK** in the Database dialog.

After you have set up row-level security in the database, you still need to set up object permissions in the repository for Presentation layer or other objects. You can also set query limits (governors). See "[Setting Up Object Permissions](#)" and "[Setting Query Limits](#)" for more information.

Setting Up Object Permissions


You can set up object permissions in your repository to control access to Presentation layer and Business Model and Mapping layer objects. You set object permissions using the Administration Tool. There are two approaches to setting object permissions: you can set permissions for particular application roles in the Identity Manager, or you can set permissions for individual objects in the Presentation layer.

This section explains how to set up object permissions for application roles in the Identity Manager. See "[Setting Permissions for Presentation Layer Objects](#)" for information about setting object permissions for individual Presentation layer objects.

Setting up object permissions for particular application roles is useful when you want to define permissions for a large set of objects at one time. You should always set up object permissions for particular application roles rather than for individual users.

[Figure 13-4](#) shows how object permissions restrict what users can see. The security rules are applied to all incoming clients and cannot be breached, even when the Logical SQL query is modified. In this example, an application role to which the Administrator belongs has been granted access to the Booked Amount column, so the Administrator can view the returned results. The user Anne Green is not a member of an application role with access to this object and cannot see the column in the Subject Area pane in Answers. Even if the request SQL is modified, results are not returned for this column because of the application role-based object permissions that have been set.

Figure 13–4 Object Permission Enforcement in the Oracle BI Server



Note the following:

- If an application role is granted or disallowed permissions on an object from multiple sources (for example, explicitly and through one or more additional application roles), the permissions are applied based on the order of precedence.
- If you explicitly deny access to an object that has child objects, users who are members of the individual application role are denied access to the child objects. For example, if you explicitly deny access to a particular logical table, you are implicitly denying access to all of the logical columns associated with that table.
- Object permissions do not apply to repository and session variables, so values in these variables are not secure. Anybody who knows or can guess the name of the variable can use it in an expression in Answers or in a Logical SQL query. Because of this, do not put sensitive data like passwords in session or repository variables.
- You can control what level of privilege is granted by default to users and application roles for repository objects without explicit permissions set. To do this, set the `DEFAULT_PRIVILEGES` parameter in the `NQSCONFIG.INI` file. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

To set up object permissions for individual application roles:

1. Open your repository in the Administration Tool.
2. Select **Manage**, then select **Identity**.
3. In the Identity Manager dialog, in the tree pane, select **BI Repository**.
4. In the right pane, select the Application Roles tab, then double-click the application role for which you want to set object permissions.

Note that if you are in offline mode, no application roles appear in the list unless you have first modified them in online mode. See "[About Applying Data Access Security in Offline Mode](#)" for more information.

5. In the Application Role dialog, click **Permissions**.

6. In the User/Application Role Permissions dialog, in the Object Permissions tab, select an object by performing one of the following steps:
 - Click the **Add** button. Then, browse to locate the object you want, select it, and then click **Select**.
 - Click the **Name** field for an empty row. Then, browse to locate the object you want, select it, and then click **Select**.
7. Assign the appropriate permission for each object. You can choose one of the following options:
 - **Read:** Only allows read access to this object.
 - **Read/Write:** Provides both read and write access to this object.
 - **No Access:** Explicitly denies all access to this object.
8. Click **OK**, then click **OK** again to return to the Identity Manager.

About Permission Inheritance for Users and Application Roles

Users can have explicitly granted permissions. They can also have permissions granted through membership in application roles, that in turn can have permissions granted through membership in other application roles, and so on. Permissions granted explicitly to a user have precedence over permissions granted through application roles, and permissions granted explicitly to the application role take precedence over any permissions granted through other application roles.

If there are multiple application roles acting on a user or application role at the same level with conflicting security attributes, the user or application role is granted the least restrictive security attribute. Any explicit permissions acting on a user take precedence over any permissions on the same objects granted to that user through application roles.


Filter definitions, however, are always inherited. For example, if User1 is a member of Role1 and Role2, and Role1 includes a filter definition but Role2 does not, the user inherits the filter definition defined in Role1.

Note that you should always define object permissions for application roles rather than for individual users.


Example 13–1 Permission Inheritance 1

You might have a user (User1) who is explicitly granted permission to read a given table (TableA). Suppose also that User1 is a member of Role1, and Role1 explicitly denies access to TableA. The resultant permission for User1 is to read TableA, as shown in [Figure 13–5](#).

Because permissions granted directly to the user take precedence over those granted through application roles, User1 has the permission to read TableA.

Figure 13–5 User Permissions and Application Role Permissions**Example 13–2 Permission Inheritance 2**

Consider the situation shown in [Figure 13–6](#).

Figure 13–6 Permissions Example

These are the resulting permissions:

- User1 is a direct member of Role1 and Role2, and is an indirect member of Role3, Role4, and Role5.
- Because Role5 is at a lower level of precedence than Role2, its denial of access to TableA is overridden by the READ permission granted through Role2. The result is that Role2 provides READ permission on TableA.
- The resultant permissions from Role1 are NO ACCESS for TableA, READ for TableB, and READ for TableC.
- Because Role1 and Role2 have the same level of precedence and because the permissions in each cancel the other out (Role1 denies access to TableA, Role2 allows access to TableA), the less restrictive level is inherited by User1. In other words, User1 has READ access to TableA.
- The total permissions granted to User1 are READ access for TableA, TableB, and TableC.

Setting Query Limits

You can manage the query environment by setting query limits (governors) in the repository for particular application roles. You can limit queries by the number of rows received, by maximum run time, and by restricting to particular time periods. You can also allow or disallow direct database requests or the Populate privilege.

You should always set query limits for particular application roles rather than for individual users.

This section contains the following topics:

- [Accessing the Query Limits Functionality in the Administration Tool](#)
- [Limiting Queries By the Number of Rows Received](#)
- [Limiting Queries By Maximum Run Time and Restricting to Particular Time Periods](#)
- [Allowing or Disallowing Direct Database Requests](#)
- [Allowing or Disallowing the Populate Privilege](#)

Accessing the Query Limits Functionality in the Administration Tool

Follow the steps in this section to access the Query Limits tab of the User/Application Role Permissions dialog.

To access the query limits functionality in the Administration Tool for a particular application role:

1. Open your repository in the Administration Tool.
2. Select **Manage**, then select **Identity**.
3. In the Identity Manager dialog, in the tree pane, select **BI Repository**.
4. In the right pane, select the Application Roles tab, then double-click the application role for which you want to set query limits.

Note that if you are in offline mode, no application roles appear in the list unless you have first modified them in online mode. See "[About Applying Data Access Security in Offline Mode](#)" for more information.

5. In the Application Role dialog, click **Permissions**.
6. In the User/Application Role Permissions dialog, click the Query Limits tab.

Limiting Queries By the Number of Rows Received

You can control runaway queries by limiting queries to a specific number of rows.

To limit queries by the number of rows received:

1. Follow the steps in "[Accessing the Query Limits Functionality in the Administration Tool](#)" to access the Query Limits tab.
2. In the **Max Rows** column, type the maximum number of rows for users to retrieve from each source database object.
3. In the **Status Max Rows** field, select one of the following options for each database:
 - **Enable:** This limits the number of rows to the value specified. If the number of rows exceeds the **Max Rows** value, the query is terminated.

- **Disable:** Disables any limits set in the **Max Rows** field.
 - **Warn:** Does not enforce limits, but logs queries that exceed the set limit in the Query log.
 - **Ignore:** Limits are inherited from the parent application role. If there is no row limit to inherit, no limit is enforced.
4. Click **OK**, then click **OK** again to return to the Identity Manager.

Limiting Queries By Maximum Run Time and Restricting to Particular Time Periods

You can forbid queries during certain time periods, or you can specify the maximum time a query can run on a database.

If you do not select a particular time period, access rights remain unchanged. If you allow or disallow access explicitly in one or more application roles, users are granted the least restrictive access for the defined time periods. For example, if a user is a member of an application role that is explicitly allowed access all day on Mondays, but that user also belongs to another application role that is disallowed access during all hours of every day, then the user has access on Mondays only.

To limit queries by maximum run time, or restrict queries to particular time periods:

1. Follow the steps in "[Accessing the Query Limits Functionality in the Administration Tool](#)" to access the Query Limits tab.
2. To specify the maximum time a query can run on a database, in the **Max Time (Minutes)** column, enter the maximum number of minutes you want queries to run on each database object. Then, in the **Status Max Time** field, select one of the following options for each database:
 - **Enable:** This limits the time to the value specified.
 - **Disable:** Disables any limits set in the **Max Time** field.
 - **Warn:** Does not enforce limits, but logs queries that exceed the set time limit in the Query log.
 - **Ignore:** Limits are inherited from the parent application role. If there is no time limit to inherit, no limit is enforced.
3. To restrict access to a database during particular time periods, in the **Restrict** column, click the **Ellipsis** button. Then, in the Restrictions dialog, perform the following steps:
 - a. To select a time period, click the start time and drag to the end time.
 - b. To explicitly grant access, click **Allow**.
 - c. To explicitly deny access, click **Disallow**.
 - d. Click **OK**.
4. Click **OK**, then click **OK** again to return to the Identity Manager.

Allowing or Disallowing Direct Database Requests

You can allow or disallow the ability to execute direct database requests for a particular application role. For the selected role, this privilege overrides the property **Allow direct database requests by default** for the database object in the Physical layer.

To set the ability to execute direct database requests:

1. Follow the steps in "[Accessing the Query Limits Functionality in the Administration Tool](#)" to access the Query Limits tab.
2. For each database object, in the **Execute Direct Database Requests** field, select one of the following options:
 - **Allow:** Explicitly grants the ability to execute direct database requests for this database.
 - **Disallow:** Explicitly denies the ability to execute direct database requests for this database.
 - **Ignore:** Limits are inherited from the parent application role. If there is no limit to inherit, then direct database requests are allowed or disallowed based on the property **Allow direct database requests by default** for the database object.
3. Click **OK**, then click **OK** again to return to the Identity Manager.

Allowing or Disallowing the Populate Privilege

When a criteria block is cached, the Populate stored procedure writes the Cache/Saved Result Set value to the database. You can grant or deny this privilege to particular application roles. For the selected application role, this privilege overrides the property **Allow populate queries by default** for the database object in the Physical layer.

Any Oracle Marketing Segmentation user who writes a cache entry or saves a result set must be a member of an application role that has been assigned the `POPULATE` privilege for the target database. For more information about marketing cache, see the topic about setting up cache for target levels in the documentation for the Oracle Marketing Segmentation application.

To allow or disallow the Populate privilege:

1. Follow the steps in "[Accessing the Query Limits Functionality in the Administration Tool](#)" to access the Query Limits tab.
2. For each database object, in the **Populate Privilege** field, select one of the following options:
 - **Allow:** Explicitly grants the Populate privilege for this database. For all Marketing data warehouses, select **Allow**.
 - **Disallow:** Explicitly denies the Populate privilege for this database.
 - **Ignore:** Limits are inherited from the parent application role. If there is no limit to inherit, then the Populate privilege is allowed or disallowed based on the property **Allow populate queries by default** for the database object.
3. Click **OK**, then click **OK** again to return to the Identity Manager.

About Applying Data Access Security in Offline Mode

It is strongly recommended that you perform data access security tasks in the Administration Tool in online mode. If you must apply data access security in offline mode, be aware that users and application roles do not appear in the Administration Tool in offline mode unless you have first modified them in the Administration Tool in online mode.

For example, if you open the Administration Tool in offline mode without first making any changes in online mode, you will see zero users and application roles defined.

However, if you first modify the users and application roles in online mode (for example, applying object permissions or setting query limits), they will subsequently be available in the Administration Tool in offline mode.

Setting Up Placeholder Application Roles for Offline Repository Development

Application roles are created and managed in the policy store using the Oracle WebLogic Administration Console and Fusion Middleware Control. These application roles are displayed in the Administration Tool in online mode so that you can use them to set data filters, object permissions, and query limits for particular roles. The application roles in the policy store are retrieved by the Oracle BI Server when it starts.

In some cases, you may want to proceed with setting up data access security in your repository for application roles that have not yet been defined in the policy store. You can do this by creating placeholder application roles in the Administration Tool, then proceeding with setting up data access security in the repository.

If you create placeholder application roles in the Administration Tool, you must eventually add them to the policy store. Run the Consistency Checker to identify application roles that have been defined in the Administration Tool, but that have not yet been added to the policy store. Be sure to use the same name in the policy store that you used for the placeholder role in the Administration Tool.

Note: Use caution when defining and using placeholder roles. If make changes to a role in offline mode that also exists in the policy store, the changes will be overwritten the next time you connect to the Oracle BI Server.

To create placeholder application roles in the Administration Tool:

1. Open your repository in the Administration Tool.
2. Select **Manage**, then select **Identity**.
3. In the Identity Manager dialog, select **Action > New > Application Role**.
4. In the Application Role dialog, provide the following information:
 - **Name:** Provide a name for the role.
 - **Description:** Optionally, provide a description of this application role.
 - **Members:** Use the **Add** and **Remove** buttons to add or remove users and other application roles as appropriate.
 - **Permissions:** Set object permissions, data filters, and query limits for this application role as appropriate. Refer to the other sections in this chapter for detailed information.
5. Click **OK** to return to the Identity Manager.

To check for application roles that need to be added to the policy store:

1. Open your repository in online mode in the Administration Tool.
2. Select **File**, then select **Check Global Consistency**.
3. Note any entries related to application roles, then add the appropriate roles to the policy store as appropriate. See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for information about adding application roles to the policy store.

4. Optionally, select individual rows and click **Copy** to copy the entries to a text file.

Alternatively, you can check an individual application role by right-clicking the application role in the Identity Manager dialog and then selecting **Check Consistency**.

About the List of Users in the Administration Tool

The Identity Manager in the Administration Tool provides a list of users that have been defined for your system. The list of users is retrieved from your authentication provider. The set of users is refreshed when the Oracle BI Server is restarted. To see the user list, select BI Repository in the Identity Manager navigation tree, and then select the Users tab in the right pane.

In online mode, by default, no users are retrieved, because the list of users might be very large. Select **Action**, then select **Set Online User Filter** to specify the set of users you want to retrieve.

The filter is empty by default, which means that no users are retrieved. Enter * to retrieve all users, or enter a combination of characters for a specific set of users, such as A* to retrieve all users whose names begin with the letter A. The filter is not case-sensitive.

In offline mode, users do not appear in the list unless you have first modified them in the Administration Tool in online mode. Because of this, you might not see any users in the Administration Tool in offline mode.

Double-click a user in the Users list to open the User dialog. You can do the following in this dialog:

- In the User tab, you can view the application roles to which this user belongs. You can also set the query logging level for this user. See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about setting the query logging level.
- In the Logons tab, you can provide a list of data source-specific logons. In this tab, you can provide a mapping of credentials that you want to be passed to data sources for this user.

This feature is used when you set up a data source connection with no shared connection pool, so that individual user names are passed directly to data sources. Rather than passing the Oracle Business Intelligence user credentials to the data source, you can map individual users to separate data source-specific credentials.

Important: Do not set object permissions, data filters, or query limits for individual users using the Permissions button. Always use application roles rather than individual users to secure data.

Completing Oracle BI Repository Setup

After you have created the repository file, the Physical layer, Business Model and Mapping layer, and Presentation layer, you need to perform several tasks to complete the initial repository setup. These tasks include saving the repository and checking consistency, adding an entry in NQSConfig.INI, and creating data source connections to the Oracle BI Server for client applications. This chapter provides information about these final setup tasks.

This chapter contains the following topics:

- [Configuring the Repository for Oracle Scorecard and Strategy Management](#)
- [Saving the Repository and Checking Consistency](#)
- [Testing and Refining the Repository](#)
- [Making the Repository Available for Queries](#)
- [Creating Data Source Connections to the Oracle BI Server for Client Applications](#)
- [Publishing to the User Community](#)

Configuring the Repository for Oracle Scorecard and Strategy Management

If your organization licensed Oracle Scorecard and Strategy Management and if you have the appropriate privileges, then you can use this functionality as part of a default installation with no additional configuration. Oracle Scorecard and Strategy Management also provides the capability to add comments (that is, annotations) or to override the status that is associated with specific dimension values for KPIs, Objectives, and Initiatives. KPI Watchlists offer the capability to add comments or to override statuses for KPIs. To enable these features, you must configure the repository to include a database object for storing the comment and status override information.

The database that you installed for use with Oracle Business Intelligence contains the Business Intelligence Platform schema, which includes required Oracle Scorecard and Strategy Management schema tables. For more information about installing a database for Oracle Business Intelligence and running the Repository Creation Assistant (RCU) to create the required schemas, see *Oracle Fusion Middleware Installation Guide for Oracle Business Intelligence*.

To configure Oracle Scorecard and Strategy Management for comments and status overrides:

1. In the Administration Tool, open the repository in online mode.

Online mode is strongly recommended for performing data access security tasks, such as the task described in Step 12 of this procedure.

2. In the Physical layer, right-click and select **New Database**. The Database dialog is displayed.
3. For **Name**, enter BSC.
4. For **Database**, select the type of database that you have installed for use with Oracle Business Intelligence (typically Oracle 11g).
5. Select the Connection Pool tab and click the **Add** button. The Connection Pool dialog is displayed.
6. For **Name**, enter BSC.
7. Select the **Call interface** appropriate for the database (for example, OCI 10g/11g for Oracle Database).
8. For **Data source name**, provide the information that is appropriate for the database that you have installed and configured for use with Oracle Business Intelligence. For example, for Oracle Database, enter a connection string similar to the following:

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=192.168.1.100)(PORT=1521))(CONNECT_
DATA=(SERVER=DEDICATED)(SERVICE_NAME=KPIOracle)(SID=KPIOracle))
```

When connecting to an Oracle Database data source, you can include the entire connect string, or you can use the net service name defined in the `tnsnames.ora` file. If you choose to enter only the net service name, then you must set up a `tnsnames.ora` file in the following location within the Oracle Business Intelligence environment, so that the Oracle BI Server can locate the entry:

```
ORACLE_HOME/network/admin
```

9. Select **Shared logon** and enter values for **User name** and **Password**. In this step, you provide the user/schema name and password that you created when you used the Repository Creation Utility (RCU) to populate the Business Intelligence Platform schema in the Oracle Business Intelligence database.

Ensure that the user that you provide has read/write privileges for the ANNOTATIONS and ASSESSMENT_OVERRIDES tables in the Business Intelligence Platform schema.
10. Click **OK** in the Connection Pool dialog.
11. Click **OK** in the Database dialog.
12. Use the Identity Manager in the Administration Tool to allow the BSystem application role to execute direct database requests by default for the BSC database object. See ["Allowing or Disallowing Direct Database Requests"](#) for more information.
13. Save and close the repository.
14. Restart the Oracle BI Server.

Saving the Repository and Checking Consistency

In offline editing, remember to save your repository from time to time. You can save a repository in offline mode even though the business models may be inconsistent.

To determine if business models are consistent, use the Check Consistency command to check for compilation errors. You can check for errors in the whole repository by choosing **File > Check Global Consistency**, or in a particular logical business model by selecting a business model and then selecting **Check Consistency** from the right-click menu.

The consistency check analyzes the repository for certain kinds of errors and inconsistencies. For example, the consistency check finds any logical tables that do not have logical sources configured or any logical columns that are not mapped to physical sources, checks for undefined logical join conditions, determines whether any physical tables referenced in a business model are not joined to the other tables referenced in the business model, and checks for existence of a subject area for each business model.

Note: Passing a consistency check does not guarantee that a business model is constructed correctly, but it does rule out many common problems.

When you check for consistency, any errors or warnings that occur are displayed in a dialog. Correct any errors and check for consistency again, repeating this process until there are no more errors. An error message indicates a problem that must be corrected. A warning message identifies a possible problem. Refer to "[Checking the Consistency of a Repository or a Business Model](#)" for more information.

After upgrading from a previous software version and checking the consistency of your repository, you might observe messages that you had not received in previous consistency checks. This typically indicates inconsistencies that had been undetected before the upgrade, not new errors.

Testing and Refining the Repository

When your repository is complete, you can run sample queries against it to test that it is created properly. Correct any problems you find and test again, repeating this process until you are satisfied with the results.

You can use the Oracle BI Server utility `nqcmd` to run test queries against the repository. The utility connects using an Oracle BI Server ODBC DSN. The Oracle BI Server must be running to use `nqcmd`.

The `nqcmd` utility is available on both Windows and UNIX systems.

This utility is intended for sanity testing. For heavier load testing, use Answers or another client. Queries with many thousands of rows will not work with `nqcmd`.

Although you can use `nqcmd` to run queries against other ODBC data sources, this section only describes how to use this utility to query the Oracle BI Server.

Before running `nqcmd`, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX) to launch a command prompt or shell window that is initialized to your Oracle instance. You can find this utility in:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup`

Then, run `nqcmd` from the resulting shell window with the desired options. For example:

```
nqcmd -dmy_dsn -umy_username [-pmy_password] -ssql_input_file -omy_result_file
```

You can pass a text file with SQL statements to the utility (script mode), or you can enter SQL at the command line (interactive mode). Queries are run against the default subject area, unless the object names used in the query are fully qualified.

Table 14–1 lists the command-line arguments for `nqcmd`.

Table 14–1 Command-line Arguments for `nqcmd`

Argument	Description
-?	Lists the available command-line arguments.
-d <i>data_source_name</i>	The ODBC data source name for the Oracle BI Server to which you want to connect. If you omit this parameter, you are prompted at the command line to enter the DSN. Tip: On Windows, you can see the available local ODBC data source names by going to Control Panel > Administrative Tools > Data Sources (ODBC) . Click the System DSN tab to see a list of the available DSNs (for example, AnalyticsWeb_coreapplication).
-u <i>user_name</i>	A valid Oracle Business Intelligence user name.
-p <i>password</i>	The corresponding Oracle Business Intelligence user password. The password argument is optional. If you do not provide a password argument, you are prompted to enter a password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide a password argument either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.
-s <i>sql_input_file_name</i>	The name and path of a text file that includes your test SQL queries.
-o <i>output_result_file_name</i>	The name and path of a file to which the utility will write the query results. This option is only used with -s.
-D <i>delimiter</i>	The delimiter used in the SQL input file (for example, semicolon (;) or colon (:)). This option is only used with -s.
-a	Enables asynchronous processing. This option is typically used with -s, when you are passing a SQL input file with multiple SQL statements.
-z	Enables UTF8 output instead of ACP in the output result file. You might need to include this option to display international characters in query results.
-utf16	Enables UTF16 instead of ACP for communication between <code>nqcmd</code> and the Oracle BI ODBC driver. You might need to include this option to display international characters in query results.
-NotForwardCursor	Disables the ODBC forward only cursor. Including this argument overrides the setting specified in the ODBC DSN.
-v	Displays the version of the <code>nqcmd</code> utility.
-SessionVar <i>session_variable_name=session_variable_value</i>	Includes the specified session variable and sets it to the specified value.

Although -C, -R, -f, -H, -q, and -NoFetch are listed by the utility as available arguments, these options are not typically used.

If you run nqcmd in interactive mode rather than script mode (or in other words, if you do not pass a SQL input file), nqcmd shows a menu of options after you provide the data source name and user credentials. Although many options are shown, you typically only use Q, T, and C against the Oracle BI Server.

Enter Q to type a query at the command line. You must enter the query on a single line, and you cannot use a semicolon as a delimiter. Pressing Enter sends the SQL to the Oracle BI Server.

Enter T to browse presentation tables, or C to browse presentation columns. The utility prompts you for catalog pattern, user pattern, table pattern, and table type pattern before returning results.

For catalog pattern, enter the subject area that contains the tables you want to see. For table pattern, enter the specific table. You can enter percent (%) to see all subject areas or all tables, use % with other characters to replace a set of characters, or use underscore (_) with other characters to replace a single character.

User pattern and table type pattern are not used in queries against the Oracle BI Server, so enter % for these options.

You can also enter D to view a static list of data types supported by the Oracle BI Server.

Making the Repository Available for Queries

After you build a repository and it is consistent, you need to upload the repository using Fusion Middleware Control so that all Oracle BI Server instances can access it. Uploading the repository allows the Oracle BI Server to load the repository into memory upon startup and makes the repository available for queries.

See "Configuring Repositories" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for information about how to perform this task.

When the repository is uploaded and you can connect to it, run sample queries against it to test that it is created properly. Correct any problems you find and test again, repeating this process until you are satisfied with the results.

Creating Data Source Connections to the Oracle BI Server for Client Applications

If you want to enable end user client applications to connect to the new repository, you must define an ODBC data source connection to the Oracle BI Server for each application. Note that Oracle BI Presentation Services has the same relationship to the Oracle BI Server as any other client application.

See "Integrating Other Clients with Oracle Business Intelligence" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for information about creating ODBC data source connections for the Oracle BI Server.

Publishing to the User Community

After testing is complete, notify the user community that the data sources are available for querying. Presentation Services users only need to know the URL to type in their

browser. Client/server users (for example, users accessing the Oracle BI Server with a query tool or report writer client application) need to know the subject area names, the computer on which the server is running, and their user IDs and passwords. They also need to have the ODBC DSN for the Oracle BI Server installed on their computers, and they may need to know the logical names of repositories if multiple repositories are used and the data source name (DSN) being created does not point to the default repository.

Setting Up Data Sources on Linux and UNIX

Most repository development is performed on Windows, because the Administration Tool runs only on Windows. When you move to a production system, however, you can choose to run the Oracle BI Server on a Linux or UNIX platform. This chapter describes how to set up data sources for use with Oracle Business Intelligence when the Oracle BI Server is running on Linux or UNIX.

See "[System Requirements and Certification](#)" for information about supported Linux and UNIX platforms.

This chapter contains the following topics:

- [About Setting Up Data Sources on Linux and UNIX](#)
- [Configuring Data Source Connections Using Native Gateways](#)
- [Using DataDirect Connect ODBC Drivers on Linux and UNIX](#)
- [Configuring Database Connections Using Native ODBC Drivers](#)
- [Configuring Oracle RPAS ODBC Data Sources on AIX UNIX](#)
- [Configuring Essbase Data Sources on Linux and UNIX](#)
- [Configuring DB2 Connect on IBM z/OS and s/390 Platforms](#)

About Setting Up Data Sources on Linux and UNIX

When the Oracle BI Server is running on Linux or UNIX, most data source connections are for query-only access. The Administration Tool is used for importing objects and is a Windows-only tool. Because of this, data source connections for import must be set up on Windows.

Note that some data source connections on Linux and UNIX do support write operations for special functions, like data source connections for write-back, usage tracking, and annotations for Oracle Scorecard and Strategy Management.

When the Oracle BI Server is running on Linux or UNIX and you need to update database object settings (such as the database type) or connection pool settings, you can copy the repository file to a Windows computer, make the changes using the Administration Tool on Windows, and then copy the repository file back to the Linux or UNIX computer.

There are three types of data source connections on Linux and UNIX platforms:

- Native data source gateway connections, such as OCI for Oracle Database or DB2 CLI for IBM DB2

- ODBC connections using the DataDirect Connect ODBC drivers that are bundled with Oracle Business Intelligence
- Native ODBC connections using external drivers, such as for Teradata data sources

Note that you cannot have a single repository that contains both DataDirect Connect ODBC connections and native ODBC connections.

Configuring Data Source Connections Using Native Gateways

You can connect to both Oracle Database and DB2 using native gateways (OCI and DB2 CLI, respectively). To configure these connections, install the appropriate database client on the computer running the Oracle BI Server, then edit the user.sh file to set environment variables for the database client.

For Oracle Database, ensure the following:

- The computer running Oracle Business Intelligence must use the Oracle Call Interface (OCI) to connect to the database.
- You must install the Oracle Database Client on the computer running the Oracle BI Server to connect to an Oracle Database.
- In the tnsnames.ora file, the Oracle Database alias (the defined entry name) must match the Data Source Name used in the repository connection pools of all physical Oracle databases.

When connecting to an Oracle Database data source, you can include the entire connect string, or you can use the net service name defined in the tnsnames.ora file. If you choose to enter only the net service name, you must set up a tnsnames.ora file in the following location within the Oracle Business Intelligence environment, so that the Oracle BI Server can locate the entry:

`ORACLE_HOME/network/admin`

Only certain versions of the Oracle Database Client are supported. See "[System Requirements and Certification](#)" for more information.

To connect to Oracle Database or DB2 using native connections:

1. Open the user.sh file. You can find this file at:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/user.sh`

2. Include the appropriate environment variable settings for the database client of your choice. Make sure to point to the appropriate libraries, depending on whether you are using a 32-bit or 64-bit database. See [Example 15–1](#) for sample values.

For an Oracle database using an OCI connection, establish the Web client before you input the tnsnames entry.

3. Save and close the file.

Example 15–1 Sample user.sh Entries for Oracle Database and DB2 (32-Bit)

This example shows sample entries in user.sh for Oracle Database and DB2 on various platforms.

```
#####
# Linux: Oracle BI 32 bit mode
#####
```

```

#set +u

# Oracle Parameters
#-----
# Make sure that Oracle DB 32 bit Client is installed
#ORACLE_HOME=/export/home/oracle/10g
#export ORACLE_HOME
#TNS_ADMIN=$ORACLE_HOME/network/admin
#export TNS_ADMIN
#PATH=$ORACLE_HOME/bin:/opt/bin:$PATH
#export PATH
#LD_LIBRARY_PATH=$ORACLE_HOME/lib:$LD_LIBRARY_PATH
#export LD_LIBRARY_PATH

# If you have Linux 64 bit Platform, and would like to run Oracle BI 32 bit
# then you must install Oracle DB 64 bit client, and this client comes with
# 32 bit libraries under $ORACLE_HOME/lib32. The LD_LIBRARY_PATH in this case
# shall be like this:
#LD_LIBRARY_PATH=$ORACLE_HOME/lib32:$LD_LIBRARY_PATH
#export LD_LIBRARY_PATH

# DB2 Parameters
#-----
#make sure the /DB2INSTANCE/sql/lib points to 32 lib file
#. /DB2INSTANCE/sql/lib/db2profile
#-----

#####
# Solaris: Oracle BI 64 bit mode
#####
#set +u

# Oracle Parameters
#-----
# Make sure to install Oracle DB 64 bit Client
#ORACLE_HOME=/export/home/oracle/10g
#export ORACLE_HOME
#TNS_ADMIN=$ORACLE_HOME/network/admin
#export TNS_ADMIN
#PATH=$ORACLE_HOME/bin:/opt/bin:$PATH
#export PATH
#LD_LIBRARY_PATH_64=$ORACLE_HOME/lib:$LD_LIBRARY_PATH_64:/opt/j2se/jre/lib/sparc
#export LD_LIBRARY_PATH_64
#-----

# DB2 Parameters
#-----
#make sure the /DB2INSTANCE/sql/lib points to 64 lib file
#. /DB2INSTANCE/sql/lib/db2profile
#LD_LIBRARY_PATH_64=/DB2INSTANCE/sql/lib/lib:$LD_LIBRARY_PATH_64
#export LD_LIBRARY_PATH_64
#-----

#####
# HPUX Itanium: Oracle BI 64 bit mode
#####
#set +u

# Oracle Parameters
#-----

```

```

#ORACLE_HOME=/export/home/oracle/10g
#export ORACLE_HOME
#TNS_ADMIN=$ORACLE_HOME/network/admin
#export TNS_ADMIN
#PATH=$ORACLE_HOME/bin:/opt/bin:$PATH
#export PATH
#SHLIB_PATH=$ORACLE_HOME/lib:$SHLIB_PATH:/opt/j2se/jre/lib/hp700
#export SHLIB_PATH
#-----

# DB2 Parameters
#-----
#make sure the /DB2INSTANCE/sql/lib points to 64 lib file
# ./DB2INSTANCE/sql/lib/db2profile
#SHLIB_PATH=/DB2INSTANCE/sql/lib:$SHLIB_PATH
#export SHLIB_PATH
#-----

#####
# AIX: Oracle BI 64 bit mode
#####
#set +u

# Oracle Parameters
#-----
#ORACLE_HOME=/export/home/oracle/10g
#export ORACLE_HOME
#TNS_ADMIN=$ORACLE_HOME/network/admin
#export TNS_ADMIN
#PATH=$ORACLE_HOME/bin:/opt/bin:$PATH
#export PATH
#LIBPATH=$ORACLE_HOME/lib:$LIBPATH:/opt/j2se/jre/lib/sparc
#export LIBPATH
#-----

# DB2 Parameters
#-----
#make sure the /DB2INSTANCE/sql/lib points to 64 lib file
# ./DB2INSTANCE/sql/lib/db2profile
#-----

```

Note that the shell script excerpts shown are examples only and are not recommendations for particular software platforms. See ["System Requirements and Certification"](#) for information about supported software platforms.

Troubleshooting OCI Connections

If you are having trouble connecting to an Oracle Database using OCI, check to ensure that the following conditions are true:

- The computer running the Oracle BI Server must use Oracle Call Interface (OCI) to connect to the database.
- If you choose not to use the entire connect string in the repository connection pool, you must ensure that a valid `tnsnames.ora` file is set up in the following location within the Oracle Business Intelligence environment, so that the Oracle BI Server can locate the entry:

`ORACLE_HOME/network/admin`

- If you choose not to use the entire connect string in the repository connection pool, ensure that the net service name in the tnsnames.ora file matches the Data Source Name used in the connection pool.

For example, in the following example of a tnsnames.ora entry, the corresponding Oracle BI repository connection pool Data Source Name is ITQA2.

```
ITQA2 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = ITQALAB2) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = ITQALAB2.corp)
    )
  )
```

The following procedure shows how to check repository database and connection pool settings against the Oracle tnsnames.ora settings.

To check that the repository database and connection pool settings are correct:

1. Open your repository in the Administration Tool.
2. In the Physical layer, double-click the database you want to check to display the Database dialog.
3. On the General tab, in the **Data source definition: Database** field, ensure that the appropriate Oracle Database version is selected. Then, click **OK**.
4. Open the Connection Pool dialog for this data source. You might need to expand the database object in the Physical layer to see the Connection Pool object.
5. In the Connection Pool dialog, check that the following is true:
 - The **Call interface** field displays the appropriate value for the release of the Oracle Database you are using.
 - The **Data source name** field displays the Oracle Database net service name that you defined in the tnsnames.ora entry.
 - The **User name** and **password** fields contain the correct values.

Change the values if necessary, then click **OK**.

6. In the Oracle Business Intelligence environment, open the tnsnames.ora file located in the following directory:


```
ORACLE_HOME/network/admin
```
7. Check that a valid net service name exists with the following characteristics:
 - Matches the connection pool settings for the Data Source Name
 - Specifies the targeted Oracle physical database

About Updating Row Counts in Native Databases

This topic applies if both of the following are true:

- You are using the Update Rowcount functionality in the Administration Tool in offline mode
- You are running a heterogeneous environment, such as the Oracle BI Server on UNIX, while remote administrators run the Administration Tool on Windows computers.

When using the Update Rowcount functionality in offline mode, the Administration Tool uses local data sources on the client computer, *not* the server data sources. Therefore, Oracle Database or DB2 clients must be configured on the Windows computer running the Administration Tool so that the following conditions are true:

- Data sources point to the same database identified in the Oracle Business Intelligence user.sh file on the UNIX server.
- The name of the local data source must also match the name of the data source defined in the Connection Pool object in the physical layer of the Oracle BI repository (.rpd) file.

If these conditions are not true, and if the server and client data sources are pointing at different databases, then erroneous updated row counts or incorrect results appear.

Using DataDirect Connect ODBC Drivers on Linux and UNIX

Oracle Business Intelligence provides DataDirect Connect ODBC drivers and driver managers for Linux and UNIX operating systems for connectivity to Microsoft SQL Server, Sybase ASE, and Informix databases.

After Oracle Business Intelligence is installed, the DataDirect Connect ODBC 32-bit drivers are installed in `ORACLE_HOME/odbc/lib`. The 64-bit drivers are installed in `ORACLE_HOME/odbc/lib64`.

Note that communication between database clients and servers is typically independent of the widths and data paths. In other words, the 32-bit database drivers can communicate with 64-bit database servers, and vice versa.

You do not need to set the ODBCINI environment variable to set up the DataDirect Connect ODBC drivers. This variable is set automatically during installation.

Refer to "[System Requirements and Certification](#)" for information about supported operating systems, databases, and driver versions for the DataDirect Connect ODBC drivers.

This section contains the following topics:

- [Configuring the DataDirect Connect ODBC Driver for Microsoft SQL Server Database](#)
- [Configuring the DataDirect Connect ODBC Driver for Sybase ASE Database](#)
- [Configuring the DataDirect Connect ODBC Driver for Informix Database](#)

Configuring the DataDirect Connect ODBC Driver for Microsoft SQL Server Database

The name of the DataDirect ODBC driver file to connect to a Microsoft SQL Server database is SEMsss23.so (SEMsss23.sl on HP-UX PA-RISC). See "[System Requirements and Certification](#)" for supported versions of Microsoft SQL Server.

To configure the DataDirect Connect ODBC Driver to connect to Microsoft SQL Server:

1. Open the user.sh file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/user.sh
```

2. In the section for your operating system, include the appropriate library path environment variable for the DataDirect Connect libraries. Make sure to point to the appropriate library, depending on whether you are using a 32-bit or 64-bit database. Note the following:

- For Solaris and Linux, the library path variable is `LD_LIBRARY_PATH`.
- For HP-UX, the library path variable is `SHLIB_PATH`.
- For AIX, the library path variable is `LIBPATH`.

For example, to set the library path variable for the 32-bit driver on Linux:

```
#####
# Linux: Oracle BI 32 bit mode
#####

#SQLServer 2000 Parameters
#-----
LD_LIBRARY_PATH=/user/local/OracleBI/odbc/lib:$ LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

3. Save and close the file.

4. Open the `odbc.ini` file. You can find this file at:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/odbc.ini`

5. Create an entry for your database, ensuring that the ODBC connection name is identical to the data source name specified in the connection pool defined in the repository. Be sure to set the Driver parameter to the file name and location of the DataDirect Connect driver for Microsoft SQL Server. In the following example, the Driver parameter is set to the 64-bit DataDirect Connect driver, and the data source name is `SQLSERVER_DB`.

```
[SQLSERVER_DB]
Driver=/usr/OracleBI/odbc/lib64/SEmsss23.so
Description=DataDirect 5.1 SQL Server Wire Protocol
Address=111.111.111.111,1433
AlternateServers=
AnsiNPW=Yes
ConnectionRetryCount=0
ConnectionRetryDelay=3
Database=dbschema_name
LoadBalancing=0
LogonID=
Password=
QuoteID=No
ReportCodePageConversionErrors=0
```

6. Save and close the `odbc.ini` file.

7. Open your repository in the Administration Tool on a Windows computer.

8. In the Physical layer, double-click the database object for the Microsoft SQL Server database.

9. Click the Features tab and scroll to the `IDENTIFIER_QUOTE_CHAR` parameter. Then, replace the value for this parameter with `'` (single quotes). The default value is double quotes (`"`).

10. Click **OK**.

11. Save and close the repository.

12. On the Linux or UNIX computer, shut down Oracle Business Intelligence.

13. Copy the repository from the Windows computer to the Linux or UNIX computer.

14. Start Oracle Business Intelligence on the Linux or UNIX computer.

Configuring the DataDirect Connect ODBC Driver for Sybase ASE Database

The name of the DataDirect ODBC driver file to connect to a Sybase ASE database is SEase23.so. See "[System Requirements and Certification](#)" for information about supported versions of Sybase ASE.

To configure the DataDirect Connect ODBC Driver to connect to Sybase ASE Database:

1. Open the user.sh file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/user.sh
```

2. In the section for your operating system, include the appropriate library path environment variable for the DataDirect Connect libraries. Make sure to point to the appropriate library, depending on whether you are using a 32-bit or 64-bit database. Note the following:

- For Solaris and Linux, the library path variable is LD_LIBRARY_PATH.
- For HP-UX, the library path variable is SHLIB_PATH.
- For AIX, the library path variable is LIBPATH.

For example, to set the library path variable for the 32-bit driver on Linux:

```
LD_LIBRARY_PATH=/user/local/OracleBI/odbc/lib:$ LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

3. Save and close the file.
4. Open the odbc.ini file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/odbc.ini
```

5. Create an entry for your database, ensuring that the ODBC connection name is identical to the data source name specified in the connection pool defined in the repository. Be sure to set the Driver parameter to the file name and location of the DataDirect Connect driver for Sybase ASE Database. For NetworkAddress, provide the IP address or fully qualified host name and the port number.

In the following example, the Driver parameter is set to the 64-bit DataDirect Connect driver, and the data source name is SybaseASE_DB.

```
[SybaseASE_DB]
Driver=/usr/OracleBI/odbc/lib64/SEase23.so
Description=DataDirect 5.3 Sybase Wire Protocol
AlternateServers=
ApplicationName=
ApplicationUsingThreads=1
ArraySize=50
AuthenticationMethod=0
Charset=
ConnectionRetryCount=0
ConnectionRetryDelay=3
CursorCacheSize=1
Database=Paint
DefaultLongDataBuffLen=1024
EnableDescribeParam=0
EnableQuotedIdentifiers=0
EncryptionMethod=0
GSSClient=native
HostNameInCertificate=
InitializationString=
```



```

Language=
LoadBalancing=0
LogonID=my_id
NetworkAddress=111.111.111.111,5005
OptimizePrepare=1
PacketSize=0
Password=
RaiseErrorPositionBehavior=0
ReportCodePageConversionErrors=0
SelectMethod=0
ServicePrincipalName=
TruncateTimeTypeFractions=0
TrustStore=
TrustStorePassword=
ValidateServerCertificate=1
WorkStationID=

```

6. Save and close the `odbc.ini` file.

Configuring the DataDirect Connect ODBC Driver for Informix Database

The name of the DataDirect ODBC driver file to connect to an Informix database is `SEifcl23.so`. See "[System Requirements and Certification](#)" for information about supported versions of Informix.

To configure the DataDirect Connect ODBC Driver to connect to Informix:

1. Open the `user.sh` file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/user.sh
```

2. In the section for your operating system, include the appropriate library path environment variable for the DataDirect Connect libraries. Make sure to point to the appropriate library, depending on whether you are using a 32-bit or 64-bit database. Note the following:
 - For Solaris and Linux, the library path variable is `LD_LIBRARY_PATH`.
 - For HP-UX, the library path variable is `SHLIB_PATH`.
 - For AIX, the library path variable is `LIBPATH`.

For example, to set the library path variable for the 32-bit driver on Linux:

```
LD_LIBRARY_PATH=/user/local/OracleBI/odbc/lib:$ LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

3. Save and close the file.
4. Open the `odbc.ini` file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/odbc.ini
```

5. Create an entry for your database, ensuring that the ODBC connection name is identical to the data source name specified in the connection pool defined in the repository. Be sure to set the `Driver` parameter to the file name and location of the DataDirect Connect driver for Informix. Also, you must specify the `HostName` parameter (you can use the fully qualified host name or the IP address) and the `PortNumber` parameter.

In the following example, the `Driver` parameter is set to the 64-bit DataDirect Connect driver, and the data source name is `Informix_DB`.

```
[Informix_DB]
Driver=/usr/OracleBI/odbc/lib64/SEifc123.so
Description=DataDirect Informix Wire Protocol
AlternateServers=
ApplicationUsingThreads=1
CancelDetectInterval=0
ConnectionRetryCount=0
ConnectionRetryDelay=3
Database=
HostName=111.111.111.111
LoadBalancing=0
LogonID=informix
Password=mypassword
PortNumber=1526
ReportCodePageConversionErrors=0
ServerName=
TrimBlankFromIndexName=1
```

6. Save and close the odbc.ini file.

Configuring Database Connections Using Native ODBC Drivers

Oracle Business Intelligence bundles UNIX ODBC drivers for some data sources, but not all. For these data sources, including Teradata, you must install your own ODBC driver, then update the user.sh and odbc.ini files to configure the data source.

To configure a database connection using a native ODBC driver:

1. Install the ODBC driver for your data source.
2. Open the user.sh file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/user.sh
```

3. In the section for your operating system, include the appropriate library path environment variable for the native ODBC driver. Make sure to point to the appropriate library, depending on whether you are using a 32-bit or 64-bit database. Note the following:
 - For Solaris and Linux, the library path variable is LD_LIBRARY_PATH.
 - For HP-UX, the library path variable is SHLIB_PATH.
 - For AIX, the library path variable is LIBPATH.

For example, to set the library path variable for Teradata on Linux:

```
LD_LIBRARY_PATH=/usr/odbc/lib:/usr/lpp/tdodbc/odbc/drivers:$ LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

Note that for Teradata in particular, you need to put /usr/odbc/lib at the beginning of the library path variable before all other entries. This step is only required for Teradata.

4. Save and close the file.
5. Open the odbc.ini file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/odbc.ini
```

6. Create an entry for your database, ensuring that the ODBC connection name is identical to the data source name specified in the connection pool defined in the repository. Be sure to set the Driver parameter to the file name and location of the

native ODBC driver for your database, with the library suffix appropriate for your operating system (for example, `.so` for Solaris and AIX, or `.sl` for HP-UX).

The following example provides details for a Teradata data source on Solaris, with a data source name of `Terav502`.

```
[Terav502]
Driver=/usr/odbc/drivers/tdata.so
Description=NCR 3600 running Teradata V2R5.2
DBCName=172.20.129.42
LastUser=
Username=
Password=
Database=
DefaultDatabase=name_of_target_database_or_user
```

Note that the `DefaultDatabase` parameter can be left empty only if you have selected the option **Require fully qualified table names** in the General tab of the Connection Pool dialog for this data source in the Administration Tool.

7. Still in the `odbc.ini` file, add an entry to the section `[ODBC Data Sources]` with the details appropriate for your data source. The following example provides details for a Teradata data source with a data source name of `Terav502`.

```
Terav502=tdata.so
```

8. Using the Administration Tool, open the repository and add the new DSN you created as the Connection Pool **Data source name** for the appropriate physical databases. See ["Creating or Changing Connection Pools"](#) for more information.
9. Restart the Oracle BI Server.

Configuring Oracle RPAS ODBC Data Sources on AIX UNIX

You can access Oracle RPAS ODBC data sources when the Oracle BI Server is running on an AIX UNIX platform. To configure this database connection, first update the `odbc.ini` file to configure the Oracle RPAS ODBC data source, then use the `rdaadmin` tool to define dimension tables as not normalized at run time.

See ["Setting Up Oracle RPAS Data Sources"](#) for information about configuring Oracle RPAS ODBC data sources on Windows.

To configure Oracle RPAS ODBC as a data source on AIX UNIX:

1. Log on as a separate telnet session.
2. Open the `odbc.ini` file. You can find this file at:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup/odbc.ini
```

3. In the RPAS data source section, edit the values. For example:

```
[RPAS Sample]
Data Source Name=RPAS Sample
Driver=[client RPASClient/lib/raix/oaodbc.so
DriverUnicodeType=1
Description=OpenRDA DSN
```

The `Data Source Name` you provide must match the value entered for `DATABASE:` in Step 3 of the following procedure. Also, you must add the line `DriverUnicodeType=1` as shown in the preceding example.

To use the rdaadmin client tool to define dimension tables as not normalized at run time:

1. Locate the rdaadmin client tool in the following location:

```
/bin/rdaadmin
```

2. Run the rdaadmin client tool by typing the following command:

```
rdaadmin
```

3. Enter appropriate text when prompted, as follows:

```
DATABASE: [Oracle_RPAS_database_name]
```

The database name must match the name given for the Data Source Name in the previous task (for example, RPAS Sample).

```
ADDRESS: [ip_address]
```

```
PORT: [port_number]
```

An example port number value is 1707.

```
CONNECT_STRING: [NORMALIZE_DIM_TABLES=NO]
```

This value treats dimension tables as not normalized at run time.

```
TYPE: []
```

```
SCHEMA_PATH: []
```

```
REMARKS: []
```

4. The RPAS environment variable OPENRDA should be declared in the Oracle BI Server session on UNIX. For example, declare the variable as follows using the 64 bit rdaadmin client tool:

```
OPENRDA_INI=/rpasclient64/config/raix/openrda.ini export OPENRDA_INI
```

Configuring Essbase Data Sources on Linux and UNIX

The Oracle BI Server uses the Essbase client libraries to connect to Essbase data sources. You must ensure that the Essbase client libraries are installed on the computer running the Oracle BI Server before you can set up a connection to Essbase data sources. See "[System Requirements and Certification](#)" for information about supported versions of the Essbase Client for use with Oracle Business Intelligence.

After you install the Essbase client libraries, ensure that the PATH environment variable includes the location of the Essbase client driver. For example, to configure access to Essbase client libraries on Linux 32-bit, add the Essbase Client Libraries directory to LD_LIBRARY_PATH:

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:Essbase_Client_Libraries_directory
```

For example:

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/export/home/Hyperion/AnalyticServicesClient/bin
export LD_LIBRARY_PATH
```

Set the appropriate variable for your operating system, using the appropriate syntax. For example, update LIBPATH on AIX, SHLIB_PATH on HP-UX Itanium, and LD_LIBRARY_PATH_64 on Solaris 64-bit.

Note: The following additional steps are required for HP-UX Itanium systems:

1. Define `ESSLANG` and `LANG`.

For example:

```
ESSLANG=English_UnitedStates.UTF-8@Binary
export ESSLANG
LANG=en_US.utf8
export LANG
```

2. Comment out `LOCALE`, `SORT_ORDER_LOCALE`, and `SORT_TYPE` in the `NQConfig.ini` file. For example:

```
[ GENERAL ]
// Localization/Internationalization parameters.
// LOCALE="English-usa";
// SORT_ORDER_LOCALE="English-usa";
// SORT_TYPE="binary";
```

You also need to ensure that an additional environment variable is set appropriately (either `ESSBASEPATH` or `ARBORPATH`, depending on your client version). For more information, see *Oracle Hyperion Enterprise Performance Management System Installation and Configuration Guide* (or the equivalent title for your client version).

Finally, you need to update `opmn.xml` with the appropriate Essbase client library information. See "[Updating Essbase Information in opmn.xml](#)" for more information.

Configuring DB2 Connect on IBM z/OS and s/390 Platforms

IBM DB2 Connect does not support the option of automatically disconnecting when an application using it receives an interrupt request.

When the native database uses DB2 Connect workstation, then you must change the setting of the parameter `INTERRUPT_ENABLED`. This parameter must be changed on any Oracle Business Intelligence computer if the database or any data source resides on IBM DB2 on a mainframe running z/OS or s/390 platforms.

Note: If IBM DB2 is used, DB2 Connect must be installed on the Oracle BI Server computer. The version of DB2 Connect must match the most recent DB2 instance that was configured as a data source.

To configure the `INTERRUPT_ENABLED` parameter:

1. Configure a database alias to be used as the native CLI Data Source Name. For example, create a new database entry using DB2 Configuration Assistant.
2. Using the database alias created and the name of the actual target DB2 database, set the `INTERRUPT_ENABLED` parameter using the following syntax:

```
uncatalog dcs db local_dcsname
catalog dcs db local_dcsname as target_dbname parms \",,INTERRUPT_ENABLED\"
```

where:

- `local_dcsname` represents the local name of the host or database (database alias name)

- `target_dbname` represents the name of database on the host or database system

Note: Be sure to use backslashes to pass the quotation marks as part of the string.

The following example uses an OS390 DB2 instance:

```
uncatalog dcs db DB2_390
catalog dcs db DB2_390 as Q10B parms "\",,INTERRUPT_ENABLED,,,,,\"
catalog database DB2_390 as DB2_390 at node NDE1EF20 authentication dcs
```

Managing Oracle BI Repository Files

This chapter provides information about topics related to managing your repository files, including comparing and merging repositories, equalizing objects, and querying and managing metadata.

This chapter contains the following topics:

- [Comparing Repositories](#)
- [Equalizing Objects](#)
- [Merging Repositories](#)
- [Querying and Managing Repository Metadata](#)
- [Changing the Repository Password](#)

Comparing Repositories

This section explains how to use the Compare repositories dialog in the Administration Tool. This feature enables you to compare all repository objects in two different repositories.

If you are using an Oracle BI Applications repository and have customized its content, you can use this option to compare your customized repository to a new version of the repository received with Oracle BI Applications.

See "[Merging Repositories](#)" for more information about merging your customized Oracle BI Applications repository with a new version of the repository.

To compare two repositories:

1. In the Administration Tool, open a repository in offline mode.
The repository that you open in this step is referred to as the current repository. See "[Using Online and Offline Repository Modes](#)" for instructions on opening a repository.
2. From the **File** menu, select **Compare**.
3. In the Select Original Repository dialog, select the repository you want to compare to the open repository.
4. In the Open Offline dialog, enter the repository password and click **OK**.
5. Use the Compare repositories dialog to review the differences between the two repositories. [Figure 16–1](#) shows the Compare repositories dialog.

Figure 16–1 Compare Repositories Dialog

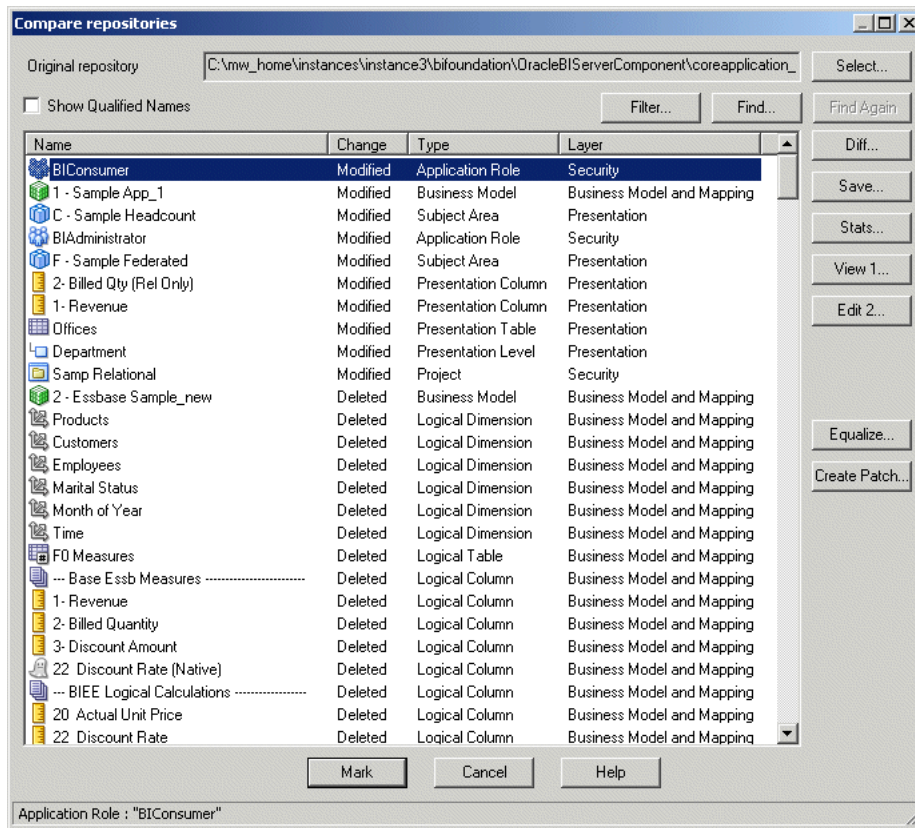


Table 16–1 lists and describes the values in the **Change** column.

Table 16–1 Compare Repositories Dialog: Change Column

Change	Description
Created	Object was created in the current repository and does not exist in the original repository.
Deleted	Object exists in the original repository but has been deleted from the current repository.
Modified	Object exists in the original repository but has been modified in the current repository.

Table 16–2 lists and describes some of the buttons in the Compare repositories dialog.

Table 16–2 Compare Repositories Dialog: Buttons

Button	Description
Filter	Opens the Comparison Filter dialog to enable you to filter the objects that appear in the Compare repositories dialog by type of change and type of object. You can specify what you want to appear and what you want to be hidden. If you select Group created and deleted objects , the tool filters out the child objects of created and deleted objects, so that only the parent objects are shown. By default, all items are shown.
Find	Search by an object Name and Type (such as Initialization Block).

Table 16–2 (Cont.) Compare Repositories Dialog: Buttons

Button	Description
Select	Enables you to select a repository to compare with the current repository.
Find Again	Search again for the most recent Find value.
Diff	Differences between the current repository and the original repository.
Save	Saves a list of the differences between the two repositories.
Stats	Provides the number of changes by Change type.
View 1	Opens an object in the original repository in read-only mode.
Edit 2	Opens an object in the current repository in read/write mode.
Equalize	Opens the Equalize Objects dialog so that you can model changes to the upgrade ID of the objects. See "Equalizing Objects" for more information.
Create Patch	Creates a patch file that contains the differences between the repositories. See "Performing Patch Merges" for more information.
Mark	Marks the object you select. Boxes appear around created and modified objects. To remove marks, from the File menu, choose Turn off Compare Mode.

Turning Off Compare Mode

This option enables you to remove marks applied to objects while using the Compare Repositories and Merge Repositories options. The **Turn off Compare Mode** option is only available after you have clicked **Mark** during the **File > Compare** action. If no repository object is marked, this option is not available.

To enable the Turn off Compare Mode option:

- In the Administration Tool, select **File**, then select **Turn off Compare Mode**.

Equalizing Objects

If you have objects in two repositories that have the same name but different upgrade IDs, you may want to treat them as the same object. To accomplish this, you can use the `equalizerpds` utility to equalize the objects by giving them both the same upgrade ID. Alternatively, you can equalize objects as part of the merge process.

You can also use the Equalize Objects dialog (available from the Compare repositories dialog) to preview what the repository will look like after you run the `equalizerpds` utility.

This section contains the following topics:

- [About Equalizing Objects](#)
- [Using the Equalize Objects Dialog](#)
- [Using the `equalizerpds` Utility](#)

About Equalizing Objects

Objects may need to be equalized because the Administration Tool tracks the history of each repository object using the upgrade ID of the object. The upgrade ID is a

unique identifier for each object. Sometimes, the upgrade ID can change because of user actions or during merge. When this occurs, and a subsequent comparison is done, the Administration Tool treats the new upgrade ID as a new object, and the missing original upgrade ID as a deleted object.

For example, assume you have two identical repositories. In one repository, delete a presentation column, then re-create it again. When you compare the two repositories using the Compare repositories dialog, there are two entries for the presentation column: one that shows the old object as deleted, and one that shows the new object as created. Without using the Compare repositories dialog, it is hard to tell that this action occurred, because the Administration Tool typically shows only the object name and properties, not the underlying upgrade ID.

It is very useful run the `equalizerpds` utility on your repositories before merging them to equalize your changes. Equalizing any opposing changes (such as a column that has been duplicated, and then renamed) cleans up the underlying upgrade IDs and prevents unintended renaming during the merge.

When you equalize objects, you can lose track of object renames because legitimate object renames become different objects. In other words, intentional renames you did in the repository might be changed to different upgrade IDs, so subsequent merges erroneously treat the renamed object as a new object. To avoid this situation, enter the before and after names of intentionally renamed objects in a rename map file that you then pass to the utility. The `equalizerpds` utility uses the information in the file to ensure that the original IDs are used in the renamed current objects.

Tip: You can view the upgrade ID for repository objects using the Query Repository dialog. To do this, follow these steps:

1. Select **Tools**, then select **Query Repository**.
 2. Run a query. See "[Querying and Managing Repository Metadata](#)" for details.
 3. Click **Columns**.
 4. Select **Upgrade ID** from the list. You can use the Find button to help locate the Upgrade ID.
 5. Click **OK**. A new column showing the upgrade IDs appears in the Results list.
-
-

Using the Equalize Objects Dialog

The Equalize Objects dialog gives you a preview of what your repository will look like if you run the `equalizerpds` utility on it. The Equalize Objects dialog provides a convenient way to compare changes related to objects that have the same name, but it does not persist any of the changes. Note that using the Equalize Objects dialog can be a very slow process for larger repositories.

To view and use the Equalize Objects dialog:

1. In the Administration Tool, open your repository in offline mode.
2. From the **File** menu, select **Compare**.
3. In the Select Original Repository dialog, select the repository you want to compare to the open repository (typically the original repository).
4. In the Open Offline dialog, enter the repository password and click **OK**. The Compare repositories dialog is displayed.
5. Click **Equalize** to display the Equalize Objects dialog.

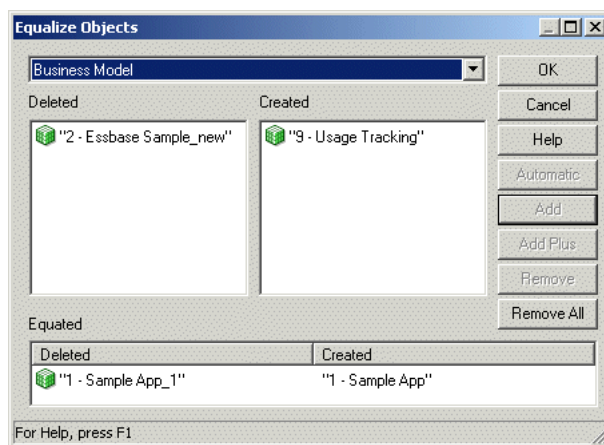
6. The Equalize Objects dialog shows a list of changes where you may want to consider objects with different upgrade IDs to be the same object. You can use the following options to model how the changes might get equalized:
 - Click **Automatic** to automatically equalize changes related to objects that have the same name. The changes appear in the Equated table.

If no changes can be automatically equalized, nothing appears in the table, and the **OK** button remains disabled.
 - Select an object in the Deleted list, then select the equivalent object in the Created list and click **Add** or **Add Plus** to equate the objects. **Add Plus** adds the object along with its child objects to the Equated table, while **Add** simply adds the selected object. For example, if you select a Subject Area and click **Add Plus**, the underlying Presentation Tables and Presentation Columns are added as well.

After you make a manual selection, the **Automatic** button is disabled.
 - Select a row in the Equated table and select **Remove** or **Remove All** to remove objects from the Equated table. **Remove All** removes the object along with its child objects, while **Remove** simply removes the selected object.

The **Automatic** button is enabled after all manual selections are removed.
7. When you are finished modeling the changes, click **OK**. The changes appear in the Compare Repositories dialog, but the changes do not persist after you close the dialog.

Figure 16–2 Equalize Objects Dialog



Using the equalizerpds Utility

You can use the `equalizerpds` utility to equalize the upgrade ID of objects in two separate repositories. If objects have the same upgrade ID, they are considered to be the same object. The utility compares upgrade IDs from the first repository (typically the original repository) with upgrade IDs from the second repository (typically the modified repository). Then, the utility equalizes the upgrade IDs of objects with the same name, using the upgrade ID from the original repository.

Before running `equalizerpds`, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX systems) to launch a command prompt that is initialized to your Oracle instance. You can find this utility in:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup`

Then, run `equalizerpds` from the resulting command prompt with the desired options. You can run the utility from this directory with no arguments or parameters to see usage information.

The utility takes the following parameters:

```
equalizerpds [-B original_repository_password] -C original_repository_name
[-E modified_repository_password] -F modified_repository_name [-J rename_map_file]
[-O output_repository_name] [-Y equalStringSet]
```

Where:

`rename_map_file` is a text file containing a list of objects that were renamed and that you want to equalize. The format is a tab-separated file with the following columns:

```
TypeName      Name1      Name2
```

For example, to include a logical column in the map file that was renamed from `Name1` to `Name2`, provide the following:

```
ATTRIBUTE      "BusinessModel"."Table"."Name1"      "BusinessModel"."Table"."Name2"
```

Do not cut and paste this example as the foundation for your own file, because the tab separators might not get copied properly. Create a new file with proper tabs.

See ["About Values for TypeName"](#) for more information about valid `TypeName` values.

`equalStringSet` is a set of characters that you want to treat as equal.

Note that the `original_repository_password` and `modified_repository_password` arguments are optional. If you do not provide these password arguments, you are prompted to enter the passwords when you run the command (`password1` and `password2`). To minimize the risk of security breaches, Oracle recommends that you do not provide password arguments either on the command line or in scripts. Note that the password arguments are supported for backward compatibility only, and will be removed in a future release.

For example:

```
equalizedrpd -C original.rpd -F modified.rpd -O modified_equalized.rpd
password1: my_original_rpd_password
password2: my_modified_rpd_password
```

In this example, `original.rpd` is compared with `modified.rpd`, the upgrade IDs are equalized using the upgrade IDs from `original.rpd`, and the final result is written to `modified_equalized.rpd`.

Note: Be sure to provide the full pathnames to your repository files, both the input files and the output file, if they are located in a different directory.

About Values for TypeName

[Table 16–3](#) shows the available object types and their corresponding values for `TypeName`.

Table 16–3 *TypeName Values*

Object Type	Value for TypeName
Database	DATABASE

Table 16-3 (Cont.) TypeName Values

Object Type	Value for TypeName
Connection Pool	CONNECTION POOL
Physical Catalog	CATALOG
Physical Schema	SCHEMA
Physical Display Folder	PHYSICAL DISPLAY FOLDER
Physical Table	TABLE
Physical Key	TABLE KEY
Physical Foreign Key	FOREIGN KEY
Physical Column	COLUMN
Physical Complex Join	JOIN
Physical Hierarchy	HIERARCHY
Physical Level	PHYSICAL LEVEL
Cube Column	COLUMN
Cube Table	CUBE TABLE
LDAP Server	LDAP SERVER
Custom Authenticator	CUSTOM AUTHENTICATOR
Variable	VARIABLE
Application Role	SECURITY ROLE
User	USER
User Database Signon	USER DATABASE SIGNON
Project	PROJECT
Business Model	SUBJECT AREA
Logical Dimension	DIMENSION
Logical Level	LEVEL
Logical Display Folder	LOGICAL DISPLAY FOLDER
Logical Table	LOGICAL TABLE
Logical Source Folder	LOGICAL SOURCE FOLDER
Logical Table Source	LOGICAL TABLE SOURCE
Logical Column	ATTRIBUTE
Logical Join	ROLE RELATIONSHIP
Logical Key	LOGICAL KEY
Logical Foreign Key	LOGICAL FOREIGN KEY
Presentation Catalog	CATALOG FOLDER
Presentation Table	ENTITY FOLDER
Presentation Column	FOLDER ATTRIBUTE
Presentation Hierarchy	PRESENTATION HIERARCHY
Presentation Level	PRESENTATION LEVEL
Catalog Link	CATALOG LINK

Table 16–3 (Cont.) TypeName Values

Object Type	Value for TypeName
Target Level	CUSTOMER TYPE
List Catalog	LIST CATALOG
Qualified Item	QUALIFIED ITEM
Qualifying Key	QUALIFYING KEY
Sampling Table	SAMPLING TABLE
Segmentation Catalog	SEGMENTATION CATALOG

Merging Repositories

You can use the Merge Repository Wizard in the Administration Tool to merge repositories (RPD files). There are three types of merges:

- Full merges are typically used during development processes, when there are two different repositories that need to be merged. The Administration Tool provides a three-way merge feature that lets you merge two repositories that have both been derived from a third, original repository. Full merges can also be used to import objects from one repository into another.
- Patch merges are used when you are applying the differential between two versions of the same repository. For example, you might want to use a patch merge to apply changes from the development version of a repository to your production repository, or to upgrade your Oracle BI Applications repository.
- Multiuser development merges are used when you are checking in projects using a multiuser development environment. See "[About the Multiuser Development Merge Process](#)" for more information.

See also [Appendix D, "Merge Rules"](#) for additional information about how repository objects are merged.

This section contains the following topics:

- [Performing Full Repository Merges](#)
- [Performing Patch Merges](#)

Performing Full Repository Merges

You can use the Administration Tool to merge different repositories. This section describes how to use the full (standard) repository merge feature in the Administration Tool.

This section contains the following topics:

- [About Full Repository Merges](#)
- [Performing Full Repository Merges With a Common Parent](#)
- [Performing Full Repository Merges Without a Common Parent](#)

About Full Repository Merges

The merge process typically involves three versions of an Oracle BI repository: the **original** repository, **modified** repository, and **current** repository. The original repository is the original unedited file (the parent repository), while the modified and

current repository are the two changed files you want to merge. The current repository is the one currently open in the Administration Tool.

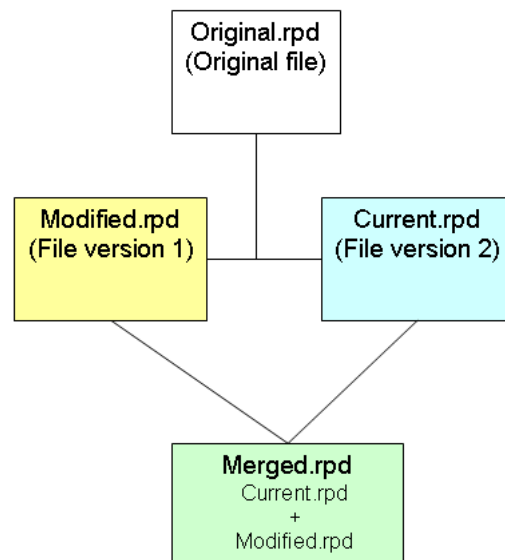
During the merge process, the Administration Tool compares the original repository with the modified repository and the original repository with the current repository. Conflicts occur when there are conflicting changes resulting from the two comparisons. For example, a conflict occurs if you rename object A to B in the modified repository, but you rename object A to C in the current repository.

The Merge Repository feature lets you decide on an object-by-object basis which changes you want to keep in the final merged repository. If there are no conflicts, merging is automatic.


There are two types of full merge:

- Common Parent.** This merge, also called a three-way merge, is useful when you have a common parent repository and two derived repositories (see [Figure 16-3](#)). There is a parent (original) RPD, and two derived RPDs (file version 1 and file version 2). After the merge, a fourth merged repository file is created.

Figure 16-3 Full Merge With a Common Parent



- No Common Parent.** This merge, also called a two-way merge, is useful when you want to import objects from one repository to another. In this case, objects are merged from two different repositories with no common parent. To accomplish this, you perform a three-way merge in the Administration Tool with a completely blank repository as the original file (see [Figure 16-4](#)). This functionality replaces the Import from Repository feature that was deprecated in an earlier release.

Figure 16–4 Full Merge Without a Common Parent

Performing Full Repository Merges With a Common Parent

This section explains how to use the Administration Tool to perform a full repository merge with a common parent. Typically, this approach is used when you have an **original** parent repository and would like to merge the changes made to objects in two modified repository versions (**current** and **modified**). Objects that do not exist in the current repository are created as new objects.

To merge two versions of an Oracle BI repository file with a common parent:

1. In the Administration Tool, open the **current** repository in offline mode.
2. From the Administration Tool menu, select **File**, then select **Merge**. The Merge Repository Wizard appears.

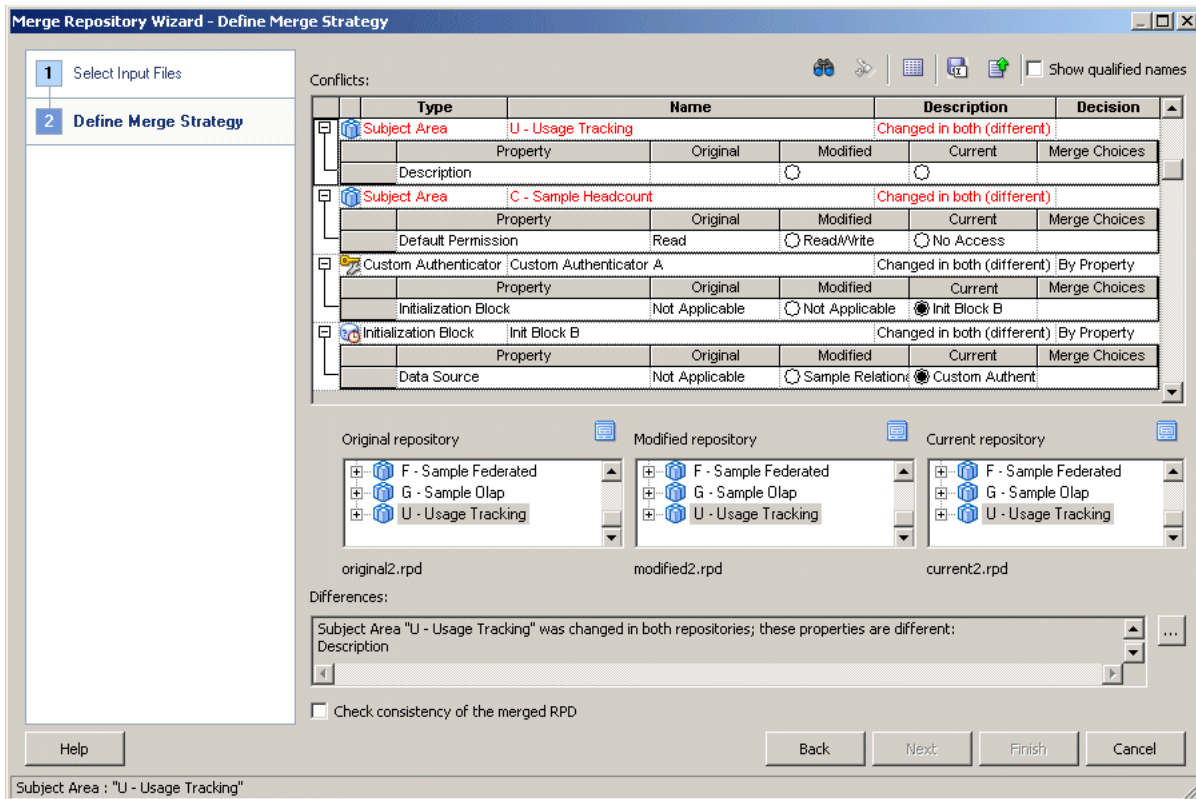
[Figure 16–5](#) shows the Merge Repository Wizard.

Figure 16–5 Merge Repository Wizard: Select Input Files Screen (Full Merge)

3. In the Select Input Files screen, for **Merge Type**, select **Full Repository Merge**.
4. Select the original parent repository by clicking **Select** next to **Original Master Repository**. Browse to select the original repository, then click **Open**.
5. Provide the password for the original repository in the appropriate **Repository Password** field.
6. Select the modified repository by clicking **Select** next to the **Modified Repository** field. Browse to select the modified repository, then click **Open**.
7. Provide the password for the modified repository in the appropriate **Repository Password** field.
8. Optionally, you can change the default name and location of the saved (merged) file by clicking **Select** next to the **Save Merged Repository as** field. Provide a new name and location, then click **Save**.
9. It is a good practice to equalize your changes to clean up underlying object IDs before merging. If you have not yet equalized your changes, select **Equalize during merge** to equalize objects as part of the merge process. Selecting this option may affect merge performance.
See "[Equalizing Objects](#)" for more information about equalizing.
10. Click **Next**. If there are any conflicts, the Define Merge Strategy screen of the Merge Repository Wizard appears. If there are no conflicts, the Merge Repository Wizard closes.

Figure 16–6 shows the Define Merge Strategy screen.

Figure 16–6 Merge Repository Wizard: Define Merge Strategy Screen



11. The Define Merge Strategy screen displays a decision table that shows conflicts for this merge. See Table 16–4 for details about the elements in this screen.

To make decisions about whether to include or exclude objects from the merged repository, choose **Current** or **Modified** from the **Decision** list. Choose **Current** to keep the change for the selected object in the current repository, or choose **Modified** to keep the change for the selected object in the modified repository.

When you select an object in the decision table, the read-only text box below the decision table describes what changes were made to that object in the current repository. In addition, the tree panels at the bottom of the dialog show the affected objects for the selected row. Alternatively, you can select an object in one of the tree views to automatically highlight the corresponding row in the decision table.

The **Modified** option in the **Decision** list displays a suffix that indicates whether the object in question will be added to or deleted from the merged repository. **Modified (A)** indicates that the object will be added, and **Modified (D)** indicates that the object will be deleted.

The type of conflict is displayed in the Description column of the Conflicts table. The decision choices you can make depend on the type of conflict shown in this column. The following list shows example results for different types of conflicts:

- Added to Current:** Choosing **Current** keeps the new object in the merged repository. Choosing **Modified (D)** deletes the new object from the merged repository.

- **Deleted from Current:** Choosing **Current** keeps the repository as it is without adding the object to the merged repository. Choosing **Modified (A)** adds the object back into the merged repository.
- **Changed in both (different):** The object was not added or deleted, but at least one of its properties was modified. Click the plus sign (+) to the left of the row to view the property that was changed, as well as its value in the original, current, and modified versions of the repository. Property values are only shown for small-length strings. Longer-length strings like descriptions, features, and init strings are not shown.

Click the option for the value you want to retain in the merged version of the repository. For some properties, such as aliases, you can choose the **Merge Choices** option to merge the properties rather than choose one over the other. This option is only available if the properties can be merged.

Note: You typically do not need to make merge decisions regarding objects that have been added to or deleted from the Modified repository. However, you can view change statistics for this merge to see a summary of changes, including objects that have been added to or deleted from Modified. See [Table 16-4](#) for more information about this feature.

After you make a merge decision, the row for that decision in the table changes from red to black. When all rows have a value in the Decision field, the **Finish** button is enabled.

12. In addition to making merge decisions, you can perform other operations in the Define Merge Strategies screen. See [Table 16-4](#) for details.
13. Click **Finish**.

[Table 16-4](#) lists and describes the elements in the Define Merge Strategies screen.

Table 16–4 Elements in the Define Merge Strategies Screen








Element	Description
Conflicts table	<p>The Conflicts table includes the following columns:</p> <ul style="list-style-type: none"> ■ Type: The type of object for which there is a conflict (for example, Presentation Column). ■ Name: The name of the object for which there is a conflict. ■ Description: The reason for the conflict, such as Added to Current. See the previous step for a description of different conflict types. ■ Decision: Select the decision according to what change you want to keep in the merged repository, such as Current, Modified (A), Modified (D), or By Property. See the previous step for a description of the results of different decisions. <p>For objects with properties that are modified in each repository, a sub-table (grid) is displayed with details of the changed properties. The grid includes the following columns:</p> <ul style="list-style-type: none"> ■ Property: The name of the property that has been modified in each repository. ■ Original: The value of the property in the original repository. ■ Modified: The value of the property in the modified repository. Select this option to keep this value. ■ Current: The value of the property in the current repository. Select this option to keep this value. ■ Merge Choices: For some properties, like aliases, you can choose this option to merge the property values rather than choose one or the other.
Show qualified names	<p>When selected, shows fully qualified names for objects in the decision table (for example, "Paint"..."Month Year Ago fact").</p> <p>Note: When the Show qualified names option is selected, some of the object names can be too long to fit into the cells of the decision table. Use the mouse to hover over the truncated name to see the full name of the object or property. Alternatively, you can manually resize columns, or double click the column separator to expand the column to the width of the object name.</p>
Check consistency of the merged RPD	<p>When selected, runs a consistency check before saving the merged file.</p>
Save Decisions to File 	<p>Saves a file containing interim changes in the Repository subdirectory so that you can stop work on the merge and continue it later. After saving the changes (decisions), close the Merge repositories dialog by clicking Cancel.</p> <p>Note: If there are a large number of decisions, you can save time by saving the merge decisions to a CSV file, opening the file in Excel or a text editor, and then modifying the merge decisions manually. Then, you can load the updated merge decisions file in the Define Merge Strategies screen.</p>
Load Decision File 	<p>Loads a saved decisions file from the Repository subdirectory so that you can continue processing a repository merge.</p>
Find by Name or Type 	<p>Searches by an object Name and Type (such as Initialization Block).</p>

Table 16–4 (Cont.) Elements in the Define Merge Strategies Screen

Element	Description
Find Again 	Searches again for the most recent Find value.
View Change Statistics 	Shows statistics for this merge, such as the number of objects deleted from the Modified repository, the number of objects that were changed in both repositories, and so on.
Details 	Shows the text in the read-only text box below the decision table in a separate window.
View Original/Modified/ Current repository 	Shows properties for the affected object in the selected repository.

Performing Full Repository Merges Without a Common Parent

This section explains how to use the Administration Tool to perform a full repository merge without a common parent. Use this method when you want to import objects from one repository (the **modified** repository) into another (the **current** repository).

Note: In the repository you choose to define as **current**, make sure that the Presentation layer references any Physical layer or Business Model and Mapping layer objects that you want to keep. Objects like business models, databases, and connection pools in the current repository that are not referenced by any Presentation layer objects are discarded during the merge. If necessary, you might want to add a placeholder subject area that references the objects before you perform the merge to ensure the desired objects are kept.

See [Appendix D, "Merge Rules"](#) for more information about which objects are retained or discarded during the merge process.

To merge two versions of an Oracle BI repository file without a common parent:

1. If you do not already have a blank repository file to serve as the original repository in the merge, create one, as follows:
 - a. In the Administration Tool, select **File**, then select **New Repository**. The Create New Repository Wizard appears.
 - b. Provide a name for the repository (for example, blank.rpd).
 - c. For **Import Metadata**, choose **No**.
 - d. Enter and confirm the repository password you want to use for this repository.
 - e. Click **Finish**.
2. Close the blank repository.
3. Open the **current** repository in offline mode. This is the repository that contains the objects you want to import.
4. From the Administration Tool menu, choose **File**, then select **Merge**. The Merge Repository Wizard appears.

5. In the Select Input Files screen, for **Merge Type**, select **Full Repository Merge**.
6. Click **Select** next to **Original Master Repository**. Then, browse to select your blank repository file as the original repository and click **Open**. Leave the password field blank.
7. Select the destination repository by clicking **Select** next to the **Modified Repository** field. Browse to select the modified repository, then click **Open**. This is the repository into which you want to import objects.
8. Provide a password for the modified repository in the appropriate **Password** field.
9. Optionally, you can change the default name and location of the saved (merged) file by clicking **Select** next to the **Save Merged Repository as** field. Provide a new name and location, then click **Save**.
10. Click **Next**. If there are any conflicts, the Define Merge Strategy screen of the Merge Repository Wizard appears. If there are no conflicts, the Merge Wizard continues with the merge process and then closes automatically when finished.
11. The Define Merge Strategy screen displays a decision table that shows conflicts for this merge. To make decisions about whether to include or exclude objects from the merged repository, choose **Current** or **Modified** from the **Decision** list. When you select an object in the decision table, the read-only text box below the decision table describes what changes were made to that object in the current repository.

Refer to [Figure 16–6](#) to see the Define Merge Strategy screen. Refer to [Table 16–4](#) for information about additional options in the Define Merge Strategy screen, such as saving merge decisions to a comma-separated values (.csv) file.

After you make a merge decision, the row for that decision in the table changes from red to black. When all rows have a value in the Decision field, the **Finish** button is enabled.
12. Click **Finish**.

Performing Patch Merges

Oracle Business Intelligence provides the capability of generating an XML patch file that contains only the changes made to a repository. This patch can be then applied to the old (original) version of the repository to create the new version. This is very useful for development-to-production scenarios, and can also be used for Oracle BI Applications customers to upgrade their repository.

This section explains how to generate a patch that contains the differences between two repositories, and then apply the patch to a repository file.

This section contains the following topics:

- [About Patch Merges](#)
- [Generating a Repository Patch](#)
- [Applying a Repository Patch](#)


About Patch Merges

In a patch merge, you create a patch that contains the differences between the **current** repository file and the **original** repository file. Then, you apply the patch file to the **modified** repository file.

In a development-to-production scenario, you have an original parent file, a current file that contains the latest development changes, and a modified file that is the deployed copy of the original file.


To generate a patch, you open the current file and select the original file, then create the patch. [Figure 16-7](#) shows how to create a patch in a development-to-production scenario.

Figure 16-7 Development-to-Production: Creating the Patch




To apply the patch, you open the modified file and select the original file, then apply the patch. [Figure 16-8](#) shows how to apply a patch in a development-to-production scenario.

Figure 16-8 Development-to-Production: Applying the Patch




In an Oracle BI Applications repository upgrade scenario, the current file is the latest version of the repository shipped by Oracle, and the original file is the original repository shipped by Oracle. The modified file is the file that contains the customizations you made to the original file.

To generate a patch, you open the current file and select the original file, then create the patch. [Figure 16-7](#) shows how to create a patch in an Oracle BI Applications repository upgrade scenario.

Figure 16–9 Oracle BI Applications Repository Upgrade: Creating the Patch

To apply the patch, you open the modified file and select the original file, then apply the patch. [Figure 16–10](#) shows how to apply a patch in an Oracle BI Applications repository upgrade scenario.

Figure 16–10 Oracle BI Applications Repository Upgrade: Applying the Patch

Generating a Repository Patch

Use the Administration Tool to generate a patch that contains the differences between two repositories.

To generate a patch using the Administration Tool:

1. In the Administration Tool, open the **current** Oracle BI repository in offline mode. In other words, open the updated repository that contains the changes you want to put in the patch.
2. Select **File**, then select **Compare**.
3. Select the **original** Oracle BI repository. When prompted, provide the appropriate password. The Compare repositories dialog appears.
4. It is a good practice to equalize your changes to clean up underlying object IDs before generating a patch. See ["Equalizing Objects"](#) for more information.
5. In the Compare repositories dialog, review the changes between the repositories. Then, click **Create Patch**.
6. In the Create Patch dialog, enter a name for the patch file (for example, my_patch.xml) and click **Save**.

Applying a Repository Patch

Use the Administration Tool to apply a patch that contains the differences between two repositories.


Note that you can apply patches from a larger multiuser repository to a smaller subset extract repository. In this case, only the changes in the subset are applied from the patch.

To apply a patch:

1. In the Administration Tool, open the **modified** Oracle BI repository in offline mode. In other words, open the repository on which you want to apply the patch.
2. Select **File**, then select **Merge**. The Merge Repository Wizard appears.

Figure 16–11 shows the Merge Repository Wizard.

Figure 16–11 Merge Repository Wizard: Select Input Files Screen (Patch Merge)



3. For **Merge Type**, select **Patch Repository Merge**.
4. Click **Select** next to **Original Master Repository**. Browse to select the original repository, then click **Open**. Note that the original repository cannot be the same as the modified (currently open) repository.
5. Enter the repository password for the original repository.
6. Click **Select** next to **Patch File**. Browse to select the patch file you want to apply, then click **Open**.
7. Optionally, click **Select** next to **Save Merged Repository as**, then enter a file name under which the patched repository will be saved and click **Save**.
8. Click **Finish**.

Using patchrpd to Apply a Patch You can also apply a patch using the patchrpd utility. This feature is especially useful when you want to patch repositories on Linux and UNIX systems where the Administration Tool is not available.

Note that unlike the Administration Tool patch feature, patchrpd does not display or resolve conflicts. If a conflict is detected, patchrpd displays a warning and exits.

Before running patchrpd, you must first run bi-init.cmd (or bi-init.sh on UNIX) to launch a command prompt or shell window that is initialized to your Oracle instance. You can find this utility in:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup
```

Then, run patchrpd from the resulting shell window with the desired options, as follows:

```
patchrpd [-P modified_rpd_password] -C modified_rpd_pathname  
[-W original_rpd_password] -G original_rpd_pathname -I xml_patch_file_pathname  
-O output_rpd_pathname [-S schema_location] [-8]
```

Where:

modified_rpd_password is the repository password for the modified repository, also called the customer or customized repository.

The password argument for the modified repository is optional. If you do not provide a password argument for the modified repository, you are prompted to enter a password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide a password argument either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.

modified_rpd_pathname is the name and location of the modified repository.

original_rpd_password is the repository password for the original repository.

The password argument for the original repository is optional. If you do not provide a password argument for the original repository, you are prompted to enter a password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide a password argument either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.

original_rpd_pathname is the name and location of the original repository.

xml_patch_file_pathname is the name and location of the XML patch file you want to apply.

output_rpd_pathname is the name and location of the RPD output file you want to generate.

schema_location is the name and location of the Oracle BI Server XML schema. If you do not specify a location, patchrpd assumes the schema file is in the default location of *ORACLE_HOME/bifoundation/server/bin/xudml1.xsd*.

-8 specifies UTF-8 encoding.

For example:

```
patchrpd -C customer.rpd -G original.rpd -I patch.xml -O patched.rpd  
Give password for customer repository: my_modified_rpd_password  
Give password for original repository: my_original_rpd_password
```

This example applies a patch called patch.xml to the customer.rpd repository, and then generates an output repository called patched.rpd.

Querying and Managing Repository Metadata

You can use repository queries to help manage repository metadata in the following ways:

- Examine and update the internal structure of the repository. For example, you can query for objects in the repository based on name, type (such as Logical Column or Presentation Hierarchy), or on a combination of name and type. You can then edit or delete objects that appear in the Results list. You can also create new objects and view parent hierarchies.
- Query a repository and view reports that show such items as all tables mapped to a logical source, all references to a particular physical column, content filters for logical sources, initialization blocks, and security and user permissions.

For example, you might want to run a report before making any physical changes in a database that might affect the repository. You can save the report to a file in comma-separated value (CSV) or tab-delimited format.

- You can save a query to run again later, or save the query results to an external file. When you save to an external file, the encoding options are ANSI, Unicode, and UTF-8.

This section contains the following topics:

- [Querying the Repository](#)
- [Querying Related Objects](#)

Querying the Repository

You can query for objects in the repository using the Query Repository tool. You can also construct a filter to filter the results, save a query, run a previously saved query, or create new repository objects.

To query a repository:

1. In the Administration Tool, open your repository.
2. Select **Tools**, then select **Query Repository**.
3. In the Query Repository dialog, complete the query information using [Table 16–5](#) as a guide.
4. Click **Query**.

[Table 16–5](#) lists the options available in the Query Repository dialog.

Table 16–5 Query Repository Options

Option	Description
Name	Use this option to search by object name. You can use an asterisk (*) wildcard character to specify any characters. The wildcard character can represent the first or last characters in the search string. Searches are not case sensitive.
Type	Select a type from the list to narrow your search to a particular type of object, or select All Types to query all objects. The list does not contain objects such as aggregate rules, logical source folders, privilege packages, and other objects that are considered internal objects.

Table 16–5 (Cont.) Query Repository Options

Option	Description
Filter	Click Filter to create or edit a filter for your query. After you create a filter, the filter criteria appear in the box to the left of the button. See " Constructing a Filter for Query Results " for more information.
Query	Click Query when you are ready to submit your query.
Save Query As	Click Save Query As to save your query to run again later. Enter the name of the saved query in the Save query as field, then click Save .
Saved Queries	Click Saved Queries to view or run previously saved queries. You can also delete saved queries. To run a previously saved query, select the row that contains the query you want to run and click Select , or double-click the row. The Saved Queries option is only available if you have previously saved queries.
Edit	After executing a query, select an object from the Results list and click Edit to edit an object in the list of query results. Not all repository objects can be edited from the results list (for example, privilege objects and user database sign-on objects). If an object cannot be edited from the results list, Edit is not available.
Delete	After executing a query, select one or more objects in the Results list and click Delete to delete the objects. After you confirm the deletion, the objects are deleted from your metadata repository.
New	Use this option to create new repository objects. First, select the type of object you want to create from the Type list, then click New . This option is not available when All Types is selected. The dialogs that appear depend on the object type that you select. For more information, refer to the sections that describe how to create that object. Note that if you choose to create a new logical dimension, you must choose whether to create a dimension with a level-based hierarchy, or a parent-child-hierarchy. Similarly, if you choose to create a new Oracle OLAP hierarchy, you must select whether you want to create a level-based or value-based hierarchy.
Show Parent	After executing a query, select an object in the Results list and click Show Parent to view the parent hierarchy of an object. If the object does not have a parent, a message appears. You cannot use Show Parent with users or application roles. In the Parent Hierarchy dialog, you can edit or delete objects. Note that if you delete an object from this dialog, any child objects of the selected object are also deleted.
Mark	After executing a query, select one or more objects in the Results list and click Mark to mark the selected objects. To unmark the objects, select them and click Mark again. Marking objects makes them easier to visually identify as you develop metadata.
Set Icon	After executing a query, select one or more objects in the Results list and click Set Icon to select a different icon for the objects. You can set special icons for objects to help visually identify them as having common characteristics. For example, you might want to pick a special icon to identify columns that will be used only by a certain user group. To change the icons back to the original icons, select the objects and click Set Icon again. Then, select Remove associated icon and click OK .
GoTo	After executing a query, select one or more objects in the Results list and click GoTo to go to the objects in the Administration Tool view of the repository. The selected objects appear highlighted in the Physical, Business Model and Mapping, or Presentation layer. Note that the Query Repository dialog closes when you choose this option.
Save	After executing a query, click Save to save query results to an external file. Then, in the Save As dialog, provide a name, file type, and encoding value for the file, then click Save .

Table 16–5 (Cont.) Query Repository Options

Option	Description
Columns	Click Columns to add additional columns of information to the results. Then, select the columns you want from the list and click OK . Note that in the Select Columns dialog, you can re-order the columns by selecting a checked column and clicking Up or Down .
Show Qualified Name	Use this option to display the fully qualified name of the objects found by the query. For example, if you query for logical columns, the default value in the Name column of the Results list is the column name. However, if you select Show Qualified Names , the value in the Name list changes to <code>businessmodel.logicaltable.column</code> .

Constructing a Filter for Query Results

Use the Query Repository Filter dialog to filter the results in the Results list of the Query Repository dialog.

The Query Repository Filter dialog contains five columns: an Item column and its operator or selection column, a Value column and its operator or selection column, and a Delete column that lets you delete the selected filter.

To construct a filter:

1. In the Administration Tool, select **Tools**, then select **Query Repository**.
2. In the Query Repository dialog, select an item in the **Results** list or select an item from the **Type** list, and then click **Filter**.
3. In the Query Repository Filter dialog, click the **Item** field. The **Item** list contains the items by which you can filter.
4. In the **Item** list, select the filter that you want to apply to the Results or Type object you selected in Step 2. Then, adjust or enter information in the **Value** column, as appropriate.

You can construct multiple filters. When you do, the **Operator** field becomes active. When the **Operator** field is active, you can set AND and OR conditions.

5. Click **OK** to return to the Query Repository dialog. The filter appears in the box to the left of the **Filter** button.

Note: If you are constructing a complex filter, you might want to click **OK** after adding each constraint to verify that the filter construction is valid for each constraint.

Example 16–1 and Example 16–2 show how to create different kinds of filters.

Example 16–1 Viewing All Databases Referenced In a Business Model

The following example shows how to create a filter that lets you view all databases referenced in a particular business model.

1. In the Query Repository dialog, select **Database** from the **Type** list, and then click **Filter**.
2. In the Query Repository Filter dialog, click the **Item** field, and then select **Related to**.

3. Click the ellipsis button to the right of the **Value** field, and in the list, choose **Select object**.
4. In the Select dialog, select the business model by which you want to filter, and then click **Select**. Your selection appears in the **Value** field.
5. Click **OK** to return to the Query Repository dialog. The filter appears in the box to the left of the **Filter** button.
6. Click **Query**. The Results list shows the databases referenced by the business model you selected.

Example 16–2 Viewing All Presentation Columns Mapped to a Logical Column

The following example shows how to create a filter that lets you view all presentation columns mapped to a particular logical column.

1. In the Query Repository dialog, select **Presentation Column** from the **Type** list, and then click **Filter**.
2. In the Query Repository Filter dialog, click the **Item** field, and then select **Column**.
3. Click the ellipsis button to the right of the **Value** field, and in the list, choose **Select object**.
4. In the Select dialog, select the column by which you want to filter, and then click **Select**. Your selection appears in the **Value** field.
5. Click **OK** to return to the Query Repository dialog. The filter appears in the box to the left of the **Filter** button.
6. Click **Query**. The Results list shows the presentation columns mapped to the logical column you selected.

Example 16–3 Nested Queries

The following example shows nested queries, where the filter itself is another query.

1. In the Query Repository dialog, select **Logical Column** from the **Type** list, and then click **Filter**.
2. In the Query Repository Filter dialog, click the **Item** field, and then select **Related to**.
3. Click the ellipsis button to the right of the **Value** field, and in the list, choose **Set Condition for Physical Column**.
4. In the new Query Repository Filter dialog, click the **Item** field, and then select **Source column**.
5. Click the ellipsis button to the right of the **Value** field, and in the list, choose **Select Object**.
6. In the Browse dialog, select a source physical column (for example, Column A) and click **Select**.
7. Click **OK** in the Query Repository Filter dialog for the subquery condition. This subquery queries all aliases for the source column you selected.
8. In the Query Repository Filter dialog for the main query, click the **Item** field in the next row and then select **Related to**.
9. Click the ellipsis button to the right of the **Value** field, and in the list, choose **Select Object**.

10. In the Browse dialog, select the same source physical column (for example, Column A) and click **Select**.
11. Select **OR** from the **Operator** list.
12. Click **OK** to return to the Query Repository dialog. The filter appears in the box to the left of the **Filter** button.
13. Click **Query**. The Results list shows a list of logical columns related to either Column A, or aliases of Column A.

Querying Related Objects

The Query Related Objects feature enables you to query objects related to one or more objects that you select from the Physical, Business Model and Mapping, or Presentation layer.

You can only use this feature with objects selected from the same layer. For example, you cannot query objects related to both a Physical layer object and a Business Model and Mapping layer object.

To query objects related to a selected object:

1. In the Administration Tool, open your repository.
2. Select one or more objects from a single layer (for example, a set of logical columns from the Business Model and Mapping layer). The objects you select must all be of the same type.
3. Right-click the objects and select **Query Related Objects**.
4. From the right-click submenu, select an object type to narrow your search to a particular type of object, or select **All Types** to query all objects related to your source objects. If you have previously made queries for this source object type, the three most recent queries are available at the top of the submenu.

After you select an object type, the Query Related Objects dialog is displayed, showing the objects related to your source objects in the **Name** list.

Table 16–6 lists the options available in the Query Repository dialog.

Table 16–6 Query Related Objects Options

Option	Description
Mark	Select one or more objects in the Name list and click Mark to mark the selected objects. To unmark the objects, select them and click Mark again. Marking objects makes them easier to visually identify as you develop metadata.
Set Icon	Select one or more objects in the Name list and click Set Icon to select a different icon for the objects. You can set special icons for objects to help visually identify them as having common characteristics. For example, you might want to pick a special icon to identify columns that will be used only by a certain user group. To change the icons back to the original icons, select the objects and click Set Icon again. Then, select Remove associated icon and click OK .

Table 16–6 (Cont.) Query Related Objects Options

Option	Description
Show Qualified Name	Use this option to display the fully qualified name of the objects found by the query. For example, if you query for logical columns, the default value in the Name list is the column name. However, if you select Show Qualified Names , the value in the Name list changes to <code>businessmodel.logicaltable.column</code> .
Show Parent	Select an object in the Name list and click Show Parent to view the parent hierarchy of an object. If the object does not have a parent, a message appears. You cannot use Show Parent with users or application roles. In the Parent Hierarchy dialog, you can edit or delete objects. Note that if you delete an object from this dialog, any child objects of the selected object are also deleted.
GoTo	Select one or more objects in the Name list and click GoTo to go to the objects in the Administration Tool view of the repository. The selected objects appear highlighted in the Physical, Business Model and Mapping, or Presentation layer. Note that the Query Related Objects dialog closes when you choose this option.

Changing the Repository Password

Each repository has a password that is used to encrypt its contents. You create the repository password when you create a new repository file, or when you upgrade a repository from a previous release of Oracle Business Intelligence.

You can change the repository password using the Administration Tool in offline mode. You cannot change the repository password when the repository is open in online mode.

After you change the repository password in the Administration Tool, you must also publish the updated repository and specify the new password in Fusion Middleware Control. Specifying the repository password in Fusion Middleware Control enables the password to be stored in an external credential store, so that the Oracle BI Server can retrieve it to load the repository.

Note: If you are using the `SampleAppLite.rpd` sample repository, you must change the default password the first time you open it in the Administration Tool, for security reasons. See "[About the SampleApp.rpd Demonstration Repository](#)" for more information about the sample repository.

To change the repository password in the Administration Tool and Fusion Middleware Control:

1. Open the repository in the Administration Tool in offline mode.
2. Select **File**, then select **Change Password**.
3. Enter the current (old) password.
4. Enter the new password and confirm it. The repository password must be longer than five characters and cannot be empty.
5. Click **OK**.

6. Save and close the repository.
7. Open a Web browser and log in to Fusion Middleware Control from the computer where the updated repository is located.
8. In the navigation tree, expand Business Intelligence and then click **coreapplication** to display the Business Intelligence Overview page.
9. Display the Repository tab of the Deployment page.
10. Click **Lock and Edit Configuration**.
11. Click **Browse** next to Repository File. Then, select the updated repository file and click **Open**.
12. Enter the new (updated) repository password in the **Repository Password** and the **Confirm Password** fields.

Make sure to specify the password that has been set in the repository. If the passwords do not match, the Oracle BI Server fails to start, and an error is logged in nqserver.log.
13. Click **Apply**, then click **Activate Changes**.
14. Return to the Business Intelligence Overview page and click **Restart**.

See "Configuring Repositories" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about the repository options in Fusion Middleware Control.

Using Expression Builder and Other Utilities

This chapter describes Expression Builder and provides instructions for creating constraints, aggregations, and other definitions within a repository. It also describes the various utilities and wizards contained in the Administration Tool.

This chapter contains the following topics:

- [Using Expression Builder](#)
- [Using Administration Tool Utilities](#)
- [Using the Calculation Wizard](#)

Using Expression Builder

You can use the Expression Builder dialogs in the Administration Tool to create constraints, aggregations, and other definitions within a repository. Expression Builder provides automatic color highlighting and other formatting enhancements to make expressions easier to build and to read.

The expressions you create with Expression Builder are similar to expressions created with SQL. Except where noted, you can use all expressions constructed with Expression Builder in SQL queries against the Oracle BI Server.

For information about using SQL with Expression Builder, and for information about the SQL functions supported by the Oracle BI Server, see [Appendix C, "Logical SQL Reference."](#)

This section contains the following topics:

- [About the Expression Builder Dialogs](#)
- [About the Expression Builder Toolbar](#)
- [About the Categories in the Category Pane](#)
- [Setting Up an Expression](#)

About the Expression Builder Dialogs

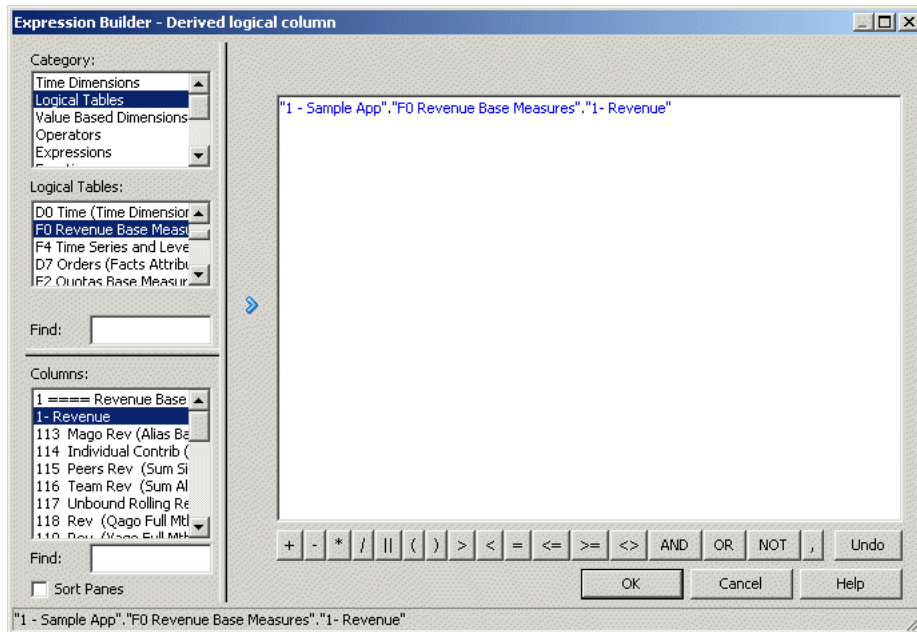
You can access Expression Builder from the following dialogs:

- Logical Table Source—Content tab
- Logical Table Source—Column Mapping tab
- Logical Column—General tab
- Logical Column—Aggregation tab

- Logical Foreign Key
- Physical Foreign Key
- Session Variable
- Repository Variable

Figure 17–1 shows Expression Builder.

Figure 17–1 Example Expression Builder Dialog



The dialog contains the following sections:

- The edit pane on the right hand side of the dialog lets you edit the current expression.
- The toolbar at the bottom of the dialog contains commonly used expression operators.
- In the left section of the dialog:
 - The top pane is the Category pane. It displays categories that are appropriate for the dialog from which you accessed Expression Builder.
 - The middle pane displays a list of available items for the category you selected in the Category pane.

You can use the **Find** field below the middle pane to display specific values in the middle pane.
 - The lower pane is the Building Blocks pane. It displays the individual building blocks for the item you selected in the middle pane.

You can use the **Find** field below the lower pane to display specific values in the lower pane.

When creating expressions in Expression Builder, you select a category from the Category pane and values are displayed in the lower panes depending on the value selected in the Category pane. When you type a value into a Find field, it filters out the non-matching strings and displays matching strings only. After typing search criteria

in a Find field, you can move up and down the list using the scroll bar, and use the tab key to move between the Find fields. To return to the full list of results, delete the string from the Find field.

Note that you can only enter text in the Find field that matches the text of one of the available strings. For example, if the available string options begin with A11, A12, and A13, the text you enter in the Find field must begin with A.

When you locate the building block you want to insert into the expression, select it and do one of the following:

- Click the arrow button
- Double click the item
- Press Enter on your keyboard

The building block you selected appears in the expression in the edit pane.

When you first open Expression Builder, the items are not sorted. When selected, the **Sort Panes** option sorts all items in the panes. As soon as you select this option, the panes are automatically redrawn without changing the contents of the panes or your filtering criteria.

About the Expression Builder Toolbar

The toolbar is located at the bottom of Expression Builder. [Table 17-1](#) describes each button and its function in an expression.

Table 17-1 Expression Builder Toolbar

Operator	Description
+	Plus sign for addition.
-	Minus sign for subtraction.
*	Multiply sign for multiplication.
/	Divide by sign for division.
	Character string concatenation.
(Open parenthesis.
)	Close parenthesis.
>	Greater than sign, indicating values higher than the comparison.
<	Less than sign, indicating values lower than the comparison.
=	Equal sign, indicating the same value.
<=	Less than or equal to sign, indicating values the same or lower than the comparison.
>=	Greater than or equal to sign, indicating values the same or higher than the comparison.
<>	Not equal to, indicating values higher or lower, but different.
AND	AND connective, indicating intersection with one or more conditions to form a compound condition.
OR	OR connective, indicating the union with one or more conditions to form a compound condition.
NOT	NOT connective, indicating a condition is not met.

Table 17–1 (Cont.) Expression Builder Toolbar

Operator	Description
,	Comma, used to separate elements in a list.

About the Categories in the Category Pane

The categories that appear in the Category pane vary, depending on the dialog from which you accessed Expression Builder. [Table 17–2](#) describes the categories that may appear.

Table 17–2 Expression Builder Categories in the Category Pane

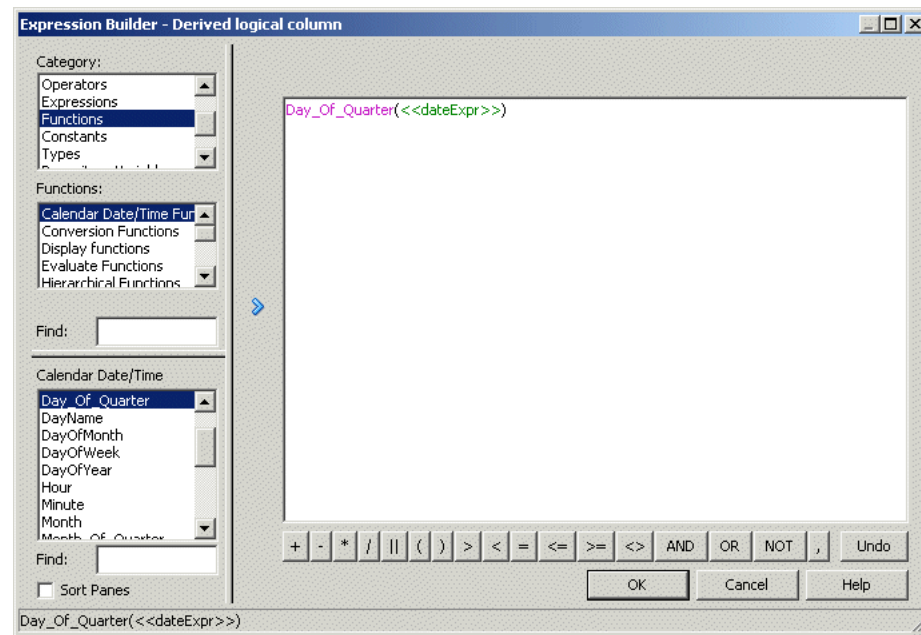
Category Name	Description
Aggregate Content	Contains the available aggregate functions. Aggregate sources must use one of the functions listed here to specify the level of their content.
Time Dimensions	Contains the time dimensions configured in the business model. If no time dimensions exist in a business model, or if time dimensions are not pertinent to a particular Expression Builder, the Time Dimensions category is not displayed. When you select the Time Dimensions category, each configured time dimension appears in the middle pane, and each level for the selected dimension appears in the lower pane.
Logical Tables	Contains the logical tables configured in the business model. If logical tables are not pertinent to a particular Expression Builder, the Logical Tables category is not displayed. When you select the Logical Tables category, each logical table in the business model appears in the middle pane, and each column for the selected logical table appears in the lower pane.
Value Based Dimensions	Contains the dimensions with parent-child hierarchies configured in the business model. If no dimensions with parent-child hierarchies exist in a business model, or if dimensions with parent-child hierarchies are not pertinent to a particular Expression Builder, the Value Based Dimensions category is not displayed. When you select the Value Based Dimensions category, the configured dimensions with parent-child hierarchies appear in the middle pane. No lower pane exists for this category.
Logical Levels	Contains the related logical levels. If level-based dimensions are not pertinent to a particular Expression Builder, the Logical Levels category is not displayed. When you select the Logical Levels category, you can then select the appropriate logical dimension (level-based) in the middle pane, and the level itself in the lower pane.
Physical Tables	Contains the related physical tables. If physical tables are not pertinent to a particular Expression Builder, the Physical Tables category is not displayed.
Operators	Contains the available SQL logical operators.
Expressions	Contains the available expressions.
Functions	Contains the available functions. The functions that appear depend on the object you selected.
Constants	Contains the available constants.
Types	Contains the available data types.

Table 17-2 (Cont.) Expression Builder Categories in the Category Pane

Category Name	Description
Repository Variables	Contains the available repository variables. If no repository variables are defined, this category does not appear.
Session Variables	Contains the available system session and non-system session variables. If no session variables are defined, this category does not appear.

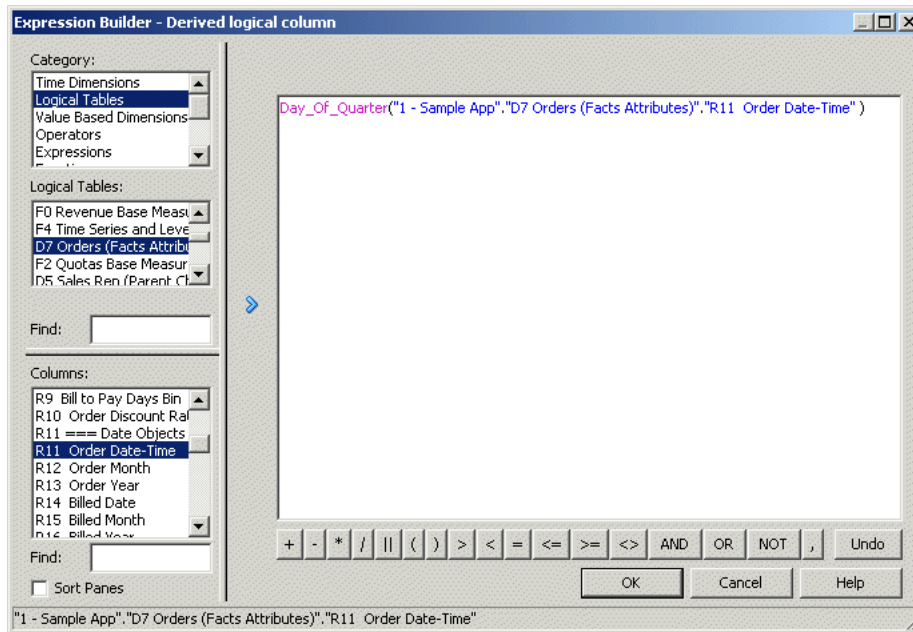
Setting Up an Expression

Figure 17-2 shows the Expression Builder dialog for a derived logical column.

Figure 17-2 Expression Builder for Derived Logical Columns

To set up an expression, select Functions from the Category pane, select a function type from Functions pane, then select a function from the lower pane. Double-click the function you want to use to paste it in the edit pane. Then, in the edit pane, click once between the parentheses of the function to select that area as the insertion point for adding the argument of the function.

To paste a logical column at the insertion point, select Logical Tables from the Category pane, select the table you want to use in the Logical Tables pane, and then double-click the logical column in the lower pane to paste the logical column at the insertion point as the argument of the function in the edit pane. Figure 17-3 shows where the expression appears in the edit pane.

Figure 17-3 Example Logical Column Function in the Editing Pane

Navigating Within Expression Builder

Use the following procedure to navigate within Expression Builder.

To navigate within Expression Builder:

1. In the Category pane, select the appropriate category for the expression you want to build.
The available expression types for the selected category appear in the middle pane.
2. Select the appropriate item for the expression you want to build.
The available building blocks for the selected item appear in the lower pane.
3. Double-click a building block to display it in the edit pane.
4. To insert an operator into the expression, click an operator on the Expression Builder toolbar.

Building an Expression

Use this procedure to build an expression in Expression Builder.

To build an expression:

1. Navigate to the individual building blocks you want in the expression.
The Syntax bar at the bottom of the Expression Builder dialog shows the syntax for the expression.
For example: BETWEEN <<Upper Bound>> AND <<Lower Bound>>
2. Add the building blocks to the edit pane.
3. Edit the building blocks to reflect the expression you want.
4. Use the Expression Builder toolbar to insert operators into the expression.
5. Repeat the preceding steps until the expression is complete, and then click **OK**.

The Administration Tool displays a message for any syntax errors in the expression. When the expression is syntactically correct, the Administration Tool adds the expression to the dialog from which you accessed Expression Builder.

Note that if the parameter `PREVENT_DIVIDE_BY_ZERO` is set to `YES` in `NQSCONFIG.INI`, the Oracle BI Server prevents errors in divide-by-zero situations, even for Answers column calculations. The Oracle BI Server creates a divide-by-zero prevention expression using `nullif()` or a similar function when it writes the physical SQL. Because of this, you do not have to use `CASE` statements to avoid divide-by-zero errors, as long as `PREVENT_DIVIDE_BY_ZERO` is set to `YES` (the default value).

See *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about configuration settings.

About the INDEXCOL Conversion Function

The `INDEXCOL` function enables you to build a derived logical column. Selecting `INDEXCOL` automatically generates the following function template:

```
IndexCol( <<integer literal>>, <<expr1>> [, <<expr2>>, ?-] )
```

Note: The argument `integer literal` can also be a session variable, an arithmetic expression, or a `CASE WHEN` statement (evaluation must be possible without reference to back-end data).

See "[INDEXCOL](#)" for more information.

Using Administration Tool Utilities

In addition to Expression Builder, the Administration Tool provides several utilities and wizards that perform functions like renaming objects, persisting aggregates, and externalizing strings.

This section contains the following topics:

- [Using the Replace Column or Table Wizard](#)
- [Using the Oracle BI Event Tables Utility](#)
- [Using the Externalize Strings Utility](#)
- [Using the Rename Wizard](#)
- [Using the Update Physical Layer Wizard](#)
- [Generating Documentation of Repository Mappings](#)
- [Generating a Metadata Dictionary](#)
- [Removing Unused Physical Objects](#)
- [Persisting Aggregates](#)

Using the Replace Column or Table Wizard

The Replace Column or Table Wizard automates the process of replacing physical columns or tables in logical table sources. For example, if you have purchased Oracle BI Applications, you can update your logical table sources to map to a different database type. You can also use this utility to change logical table source mappings from a development table to a production table.

You can use the Replace Column or Table Wizard to replace a single column (within the same table), or to replace an entire table. If you replace a table, you must map all the columns in the table.

To replace a physical column in logical table sources:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Replace Column or Table in Logical Table Sources** and click **Execute**.
3. In the Select Object screen, select **Replace single column only**.
4. In the left pane, select the physical column that you want to replace. You must select a valid column. For example, you cannot select a column that is used in a logical table source that has more than one table as a source.
5. In the right pane, select the physical column that you want to use as a replacement for the original column. Then, click **Next**.
6. The Select Sources screen shows all logical table sources that map to the physical column you selected. Select the logical table sources in which you want to change the physical column mapping. Select **Show Qualified Names** to see the full context for each source.

If you select an invalid logical table source, or in other words, one that cannot be used for replacement, a message appears explaining why that source cannot be used, and the check box for that source is disabled.

Note that invalid logical table sources do not appear in the list when **Hide unusable logical table sources in Replace wizard** has been selected in the General tab of the Options dialog. Instead, the **Info** button is displayed when a logical table source that maps to that column does not appear in the list. Click **Info** to see details on why the physical objects could not be replaced in the logical table source or sources.

The Select Sources screen only appears if there are multiple logical table sources that map to the physical column you selected.

Click **Next** after you have selected logical table sources.

7. When the repository is open in online mode, the Checkout screen appears. In online mode, objects need to be checked out before you can make changes to them. Click **Next** to check out the necessary objects.
8. The Finish screen displays a summary of the objects that will be replaced. If you want to make changes, click **Back**, or select a particular step from the navigation panel.
9. Click **Finish**.

To replace a physical table in logical table sources:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Replace Column or Table in Logical Table Sources** and click **Execute**.
3. In the Select Object screen, select **Replace whole table**.
4. In the left pane, select the physical table that you want to replace.
5. In the right pane, select the physical table that you want to use as a replacement for the original table. Then, click **Next**.
6. The Select Sources screen shows all logical table sources that map to the physical table you selected. Select the logical table sources in which you want to change the

physical table mapping. Select **Show Qualified Names** to see the full context for each source.

If you select an invalid logical table source, or in other words, one that cannot be used for replacement, a message appears explaining why that source cannot be used, and the check box for that source is disabled.

Note that invalid logical table sources do not appear in the list when **Hide unusable logical table sources in Replace wizard** has been selected in the General tab of the Options dialog. Instead, the **Info** button is displayed when a logical table source that maps to that column does not appear in the list. Click **Info** to see details on why the physical objects could not be replaced in the logical table source or sources.

The Select Sources screen only appears if there are multiple logical table sources that map to the physical table you selected.

Click **Next** after you have selected logical table sources.


7. The bottom pane of the Select Columns screen shows individual column mappings between the selected physical tables. If column names in the two selected tables match, default column mappings appear in the bottom pane.

To add a column mapping to the list of mapped columns, first select a source column in the left pane. Then, select a replacement column in the right pane and click **Add**.

To remove a column mapping from the list of mapped columns, select a row of mapped columns from the list and click **Remove**.

Figure 17–4 shows the Select Columns screen.

Figure 17–4 Select Columns Screen of the Replace Column or Table Wizard



8. When you have finished mapping columns between the selected physical tables, click **Next**.
9. When the repository is open in online mode, the Checkout screen appears. In online mode, objects need to be checked out before you can make changes to them. Click **Next** to check out the necessary objects.

10. The Finish screen displays a summary of the objects that will be replaced. If you want to make changes, click **Back**, or select a particular step from the navigation panel.
11. Click **Finish**.

Using the Oracle BI Event Tables Utility

This utility lets you identify a table as an Oracle BI event polling table. An event polling table is a way to notify the Oracle BI Server that one or more physical tables have been updated. Each row that is added to an event table describes a single update event. The cache system reads rows from, or polls, the event table, extracts the physical table information from the rows, and purges cache entries that reference those physical tables.

For more information about event polling tables, see "Cache Event Processing with an Event Polling Table" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

To start the Oracle BI Event Tables utility:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Oracle BI Event Tables** and click **Execute**.

Using the Externalize Strings Utility

You can use the Externalize Strings utility to localize the names of Presentation layer subject areas, tables, hierarchies, columns, and their descriptions. You can save these text strings to an external file with ANSI, Unicode, and UTF-8 encoding options.

Before you can use the Externalize Strings utility, you must externalize strings in the Presentation layer. Note the following about externalizing strings in the Presentation layer:

- You can right-click any Presentation layer object, such as a subject area, presentation table, or presentation column, and choose **Externalize Display Names > Generate Custom Names** or **Externalize Descriptions > Generate Custom Descriptions** to externalize strings.
- Choosing one of these right-click externalization options automatically selects the **Custom display name** or **Custom description** options in the Properties dialog for the selected object and all of its child objects.

For example, if you right-click a subject area and choose one of the externalization options, the externalization flag is set on all presentation tables, columns, hierarchies, and levels within that subject area.

- Running the Externalize Strings utility only externalizes those strings that have been selected for externalization in the Presentation layer.

For full information about using the Externalize Strings utility, see "Localizing Metadata Names in the Repository" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

To start the Externalize Strings utility:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Externalize Strings** and click **Execute**.

Using the Rename Wizard

You can use the Rename Wizard to rename tables and columns in the Presentation layer and Business Model and Mapping layer. It provides a convenient way to transform physical names to user-friendly names.

It is a best practice to rename objects in the Business Model and Mapping layer rather than the Presentation layer, for better maintainability. Giving user-friendly names to logical objects rather than presentation objects ensures that the names can be reused in multiple subject areas. Also, it ensures that the names persist even when you need to delete and re-create subject areas to incorporate changes to your business model.

Be aware that when you use the Rename Wizard to rename presentation columns, the **Use Logical Column Name** property gets set to false.

To use the Rename Wizard:

1. In the Administration Tool, select **Tools**, then select **Utilities**. Then, select **Rename Wizard** and click **Execute**.

You can also access the Rename Wizard by right-clicking an object or set of objects in the Business Model and Mapping layer or Presentation layer, and then selecting **Rename Wizard**. The wizard starts in the Select Rules screen and only applies to the logical or presentation objects you selected.

2. In the Select Objects screen, select the objects you want to rename. First, select the layer that contains the objects (**Presentation** or **Business Model and Mapping**), then select an object and click **Add**. Click **Add Hierarchy** to add all objects associated with the selected object.

Click **Next** after you have selected the objects you want to rename.

3. In the Select Types screen, select the object types you want to rename, such as Subject Area, Logical Table, or Logical Column. Then, click **Next**.
4. In the Select Rules screen, select the renaming rules you want to apply and click **Add**. Select **Change specified text** to rename particular words or phrases.

The renaming rules are applied in the order in which they appear in the list. Select a rule that you have added and click **Up** or **Down** to change the order in which the rules will be applied.


For example, say you want to rename the logical columns GlobalGROUP, GlobalSales, and GlobalCustomerName to Group, Sales, and Customer Name. To achieve this, you can apply the following rules in the given order:

```
Insert space before each first uppercase letter, unless on the first position
or there is a space already
All text lowercase
First letter of each word capital
Change each occurrence of "Global " to "" (not case sensitive)
```

Click **Next** after you have selected renaming rules.

5. When the repository is open in online mode, the Checkout screen appears. In online mode, objects need to be checked out before you can make changes to them. Click **Next** to check out the necessary objects.
6. The Finish screen displays a summary of the objects that will be renamed. If you want to make changes to the list of renamed objects, click **Back**, or select a particular step from the navigation panel. Click **Finish** to rename the objects.

Figure 17–5 shows the Rename Wizard.

Figure 17–5 Rename Wizard

Using the Update Physical Layer Wizard

You can use the Update Physical Layer Wizard to update database objects in the Physical layer of a repository, based on their current definitions in the back-end database. This wizard is not available for repositories that are open in read-only mode, because they are not available for updating.

When the wizard processes the update, the Oracle BI Server connects to each back-end database. The objects in the Physical layer are compared with those in the back-end database. Explanatory text alerts you to differences between objects as defined in the database in the Physical layer and as defined the back-end database, such as data type-length mismatches and objects that are no longer found in the back-end database. For example, if an object exists in the database in the Physical layer of the repository but not in the back-end database, the following text is displayed:

```
Object does not exist in the database and will be deleted
```

The wizard does not add columns or tables to the repository that exist in the back-end database, but not in the repository. Additionally, the wizard does not update column key assignments. It checks that there is a column in the repository that matches the column in the database, and then, if the values do not match, the wizard updates the type and length of the column in the repository.

The connection pool settings for each database need to match the connection pool settings used when the objects were last imported into the Physical layer from the back-end database. See "[Creating or Changing Connection Pools](#)" for more information about connection pool settings.

To update objects in the Physical layer:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Update Physical Layer** and click **Execute**.

The databases in the Physical layer of the repository are listed in the left pane of the wizard.

3. In the Select Database screen, select the databases that you want to update in the left pane, and then click **Add**. To remove a database from the update list in the right pane, select it and click **Remove**.
4. Click **Next**.
5. In the Select Connection Pool screen, select the connection pool for each database that you want to update and then click **Next**. You might need to set or confirm values for variables before continuing.
6. In the Update screen, review the information about each update and select the updates you want to perform. You can sort the rows (toggle between ascending and descending order) by clicking the **Name** column heading.
7. If you decide that you do not want the wizard to update a particular object in the Physical layer, click the **Back** button and remove the object.
8. When the repository is open in online mode, the Checkout screen appears. In online mode, objects need to be checked out before you can make changes to them. Click **Next** to check out the necessary objects.
9. Click **Finish**.
The wizard updates the objects in the Physical layer, and then closes automatically. Objects that do not exist in the database are deleted.
10. From the **File** menu, select **Save** to save the updated objects in the Physical layer.

Generating Documentation of Repository Mappings

The Repository Documentation utility documents the mapping from the presentation columns to the corresponding logical and physical columns. The documentation also includes conditional expressions associated with the columns. The documentation can be saved in comma separated (CSV), XML, or tab delimited (TXT) format.

You can use the Repository Documentation utility to extract Oracle Business Intelligence metadata to a flat file so that it can be loaded into Excel and RDBMS. You can query the resulting file to answer questions such as "If I delete physical column X, what logical columns will be affected?" or "How many places in the business model refer to the physical table W_SRVREQ_F?" Then, you can establish dependency relationships among elements in the repository.

Excel only allows data sets of 1,000,000 rows. You might exceed this in a large repository. Therefore, you might want to run the Repository Documentation utility on a subset of the repository by extracting relevant business models into a new project. For more information, see [Chapter 3](#).

The Repository Documentation utility creates a comma-separated values file or a tab-separated values file that shows the connections between the Presentation and Physical layers in the current repository. This file can be imported into a repository as a Physical layer. Note that the file does not include information about repository variables and marketing objects.

To run the Repository Documentation utility:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Repository Documentation** and click **Execute**.
3. In the Save As dialog, choose the directory where you want to save the file.
4. Type a name for the file.

5. Choose a type of file and an Encoding value, and then click **Save**. Current encoding options are ANSI, Unicode, and UTF-8.

Generating a Metadata Dictionary

You can generate a metadata dictionary to help Oracle Business Intelligence users obtain more information about metrics or attributes for repository objects. For example, users might need to resolve issues caused by confusing metadata object names, or to obtain more details when an attribute is derived in a complicated way.

A metadata dictionary is a static set of XML documents. Each XML document describes a metadata object, such as a column, including its properties and relationships with other metadata objects. These XML documents can be viewed within the Oracle BI Presentation Services user interface, or they can be viewed directly in a browser.

Use the Administration Tool to generate a metadata dictionary for your repository. Because the dictionary does not change dynamically as repository changes are made, you must generate the dictionary periodically to update the content.

The metadata dictionary files need to be hosted on a Web server, such as Oracle HTTP Server or Apache HTTP Server. When you generate the dictionary, you can set the output location to the final location on the Web server, or to a temporary location. If you generate the dictionary in a temporary location, you must then copy the files to the location on the Web server.

Note that some large repositories can contain tens of thousands of objects. Generating a dictionary for a large repository can take a significant period of time.

To generate a metadata dictionary:

1. In the Administration Tool, open your repository in offline mode. You cannot generate a metadata dictionary in online mode.
2. Select **Tools**, then select **Utilities**.
3. Select **Generate Metadata Dictionary** and click **Execute**.
4. In the Choose Directory dialog, click **Browse** to locate and select the location where you want to store the dictionary. You can select a destination for your dictionary in the following ways:
 - Select a local or network location. When the dictionary is generated, a subdirectory with the same name as the repository is created in that location. The dictionary directories and files are created in that subdirectory.

For example, if you select J:\BI_DataDictionary and your repository name is demo1.rpd, the dictionary files, including the style sheets, will be located in J:\BI_DataDictionary\demo1.
 - If you want to use an IIS virtual directory, you can create or select a virtual directory in IIS before you generate the dictionary. When you generate the dictionary, choose the physical directory associated with the IIS virtual directory.
5. Click **OK**.
6. If you did not save the files directly to a location on a Web server, copy the files over to your Web server and ensure they are accessible. Refer to the documentation for your Web server for detailed information.

The location where the metadata dictionary files can be viewed is dependent on the host name and port number of your Web server, along with the directory location where you store the files.

7. You must edit the `instanceconfig.xml` configuration file to enable the metadata dictionary feature in the Oracle BI Presentation Services user interface, as well as grant the appropriate privilege to your users, groups, or application roles. See "Providing Access to Metadata Dictionary Information" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information about these additional configuration steps.

After you generate a metadata dictionary, style sheets and index files are created for that dictionary. The related style sheets (XSL files) are created and stored in a directory named `xsl` within the repository directory.

A name index and tree index are created and stored in the `[drive]:\[path]\[repository name]` root directory. The index files are associated with each other so that you can quickly switch views.

For additional information about viewing metadata dictionary information from the Oracle BI Presentation Services user interface, see "Viewing Metadata Information from the Subject Areas Pane" in *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition*.

Removing Unused Physical Objects

Large repositories use more memory on the server and are harder to maintain. Additionally, development activities take longer on a large repository. This utility enables you to remove objects that you no longer need in your repository. You can remove databases, initialization blocks, physical catalogs, and variables.

To remove unused physical objects:

1. In the Administration Tool, select **Tools**, then select **Utilities**.
2. Select **Remove Unused Physical Objects** and click **Execute**.
3. In the Remove Unused Physical Objects dialog, from the **Type** list, select the type of object.
4. In the list of objects, verify that only the objects that you want to remove are selected.

Below the list of objects, the number of selected objects and the total number of objects appears.

5. To remove the selected objects, click **Yes**.

Persisting Aggregates

You can use the Aggregate Persistence Wizard to create the SQL file that will be used to create aggregate tables and map them into the metadata. See ["Using the Aggregate Persistence Wizard to Generate the Aggregate Specification"](#) for more information.

Using the Calculation Wizard

You can use the Calculation Wizard to create new calculation columns that compare two existing columns, and also to create metrics in bulk. It has a built-in mechanism to handle divide-by-zero and null cases, as well as other difficult situations. The

Calculation Wizard provides an automated way to calculate the sales by quarter, the percentage of revenue, minimum and maximum values, and so on.

To start the Calculation Wizard, right-click any logical fact or dimension column in the Business Model and Mapping layer of data type numeric, and then select the option **Calculation Wizard**. The wizard starts with the column on which you right-clicked as the source column, and then displays the other columns in that table for comparison.

To use the Calculation Wizard:

1. Right-click a measure column in the Business Model and Mapping layer (any logical fact or dimension column of data type numeric), and then select **Calculation Wizard**.
2. The first time you use the Calculation Wizard, the Introduction screen appears. Select **In the future, do not show this introduction screen** if you do not want this screen to display subsequently. If you choose not to display the Introduction screen, you can go to **Tools > Options** to cause it to appear again. See "[Setting Preferences](#)" for more information.

Click **Next** to display the Select Columns screen.

3. Select the columns that you want to compare with the source column. If the source column is mapped to multiple logical tables, a list of tables appears in the upper left pane. Select a table, then select a column or columns from the upper right pane to add comparison columns to the Selected Columns list.

You can remove items from the Selected Columns list by selecting a column and clicking **Remove**.

Click **Next** when you have finished selecting comparison columns.

4. In the New Calculations screen, you can choose which calculations you want to perform for specific columns. The elements of the New Calculation screen are as follows:
 - The upper left pane shows the name of the source column, followed by a list of comparison columns that you selected in the Select Columns screen. Select a particular column to create calculations for that column.
 - The upper right pane shows a list of calculations you can perform. Select a calculation to view the calculation definition and the default calculation name. In the calculation definition, **CurrentX** refers to the value of the source column, and **ComparisonX** refers to the value of the comparison column you selected in the upper left pane.

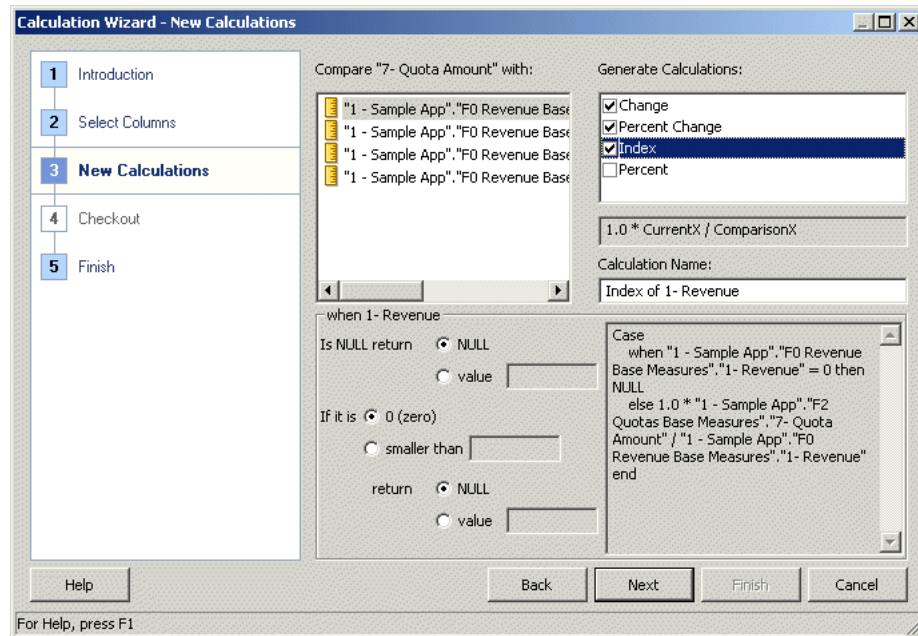
You can optionally change the calculation name. This name becomes the name of the resulting calculation column.

The following calculations are available:

- **Change (CurrentX - ComparisonX)**. Subtract the value of the comparison column from the source column.
- **Percent Change (100.0 * (CurrentX - ComparisonX) / ComparisonX)**. Subtract the value of the comparison column from the source column and express as a percentage.
- **Index (1.0 * CurrentX / ComparisonX)**. Divide the source column by the comparison column.
- **Percent (100.0 * (CurrentX / ComparisonX))**. Divide the source column by the comparison column and express as a percentage.

- The lower left pane shows special cases that are available for the selected calculation. You can keep the default values, or specify how you want the special cases to be handled. For example, for the Change calculation, you can choose whether to return NULL or some other value when the comparison column is NULL.
- Select a calculation in the upper right pane to view and set special cases for that calculation.
- The lower right pane shows the resulting SQL for the selected calculation.

Figure 17–6 *New Calculations Screen of Calculation Wizard*



5. Click **Next** when you have finished creating calculations.
6. When the repository is open in online mode, the Checkout screen appears. In online mode, objects need to be checked out before you can make changes to them. Click **Next** to check out the necessary objects.
7. The Finish screen displays a summary of the calculations that will be created. If you want to make changes, click **Back**, or select a particular step from the navigation panel.
8. Click **Finish**. The new calculation columns are created.

Using Variables in the Oracle BI Repository

You can use variables in a repository to streamline administrative tasks and dynamically modify metadata content to adjust to a changing data environment. There are two classes of variables: repository variables and session variables.

- A repository variable has a single value at any point in time. There are two types of repository variables: static and dynamic.
- Session variables are created and assigned a value when each user logs on. There are two types of session variables: system and nonsystem.

Initialization blocks are used to initialize dynamic repository variables, system session variables, and nonsystem session variables.

You can use the Variable Manager in the Administration Tool to define variables. The Variable Manager dialog has two panes. The left pane displays a tree that shows variables and initialization blocks, and the right pane displays details of the item you select in the left pane. Repository variables and system and nonsystem session variables are represented by a question mark icon. The icon for an initialization block is a cube labeled i.

Caution: Values in repository and session variables are not secure, because object permissions do not apply to variables. Anybody who knows or can guess the name of the variable can use it in an expression in Answers or in a Logical SQL query. Because of this, do not put sensitive data like passwords in session or repository variables.

This chapter contains the following topics:

- [About Repository Variables](#)
- [Creating Repository Variables](#)
- [About Session Variables](#)
- [Creating Session Variables](#)
- [Working with Initialization Blocks](#)

About Repository Variables

A repository variable has a single value at any point in time. Repository variables can be used instead of literals or constants in Expression Builder in the Administration Tool. The Oracle BI Server substitutes the value of the repository variable for the variable itself in the metadata.

This section contains the following topics:

- [About Static Repository Variables](#)
- [About Dynamic Repository Variables](#)

About Static Repository Variables

The value of a static repository variable is initialized in the Variable dialog. This value persists, and does not change until an administrator decides to change it.

For example, suppose you want to create an expression to group times of day into different day segments. If Prime Time were one of those segments and corresponded to the hours between 5:00 PM and 10:00 PM, you could create a CASE statement like the following:

```
CASE WHEN "Hour" >= 17 AND "Hour" < 23 THEN 'Prime Time' WHEN... ELSE...END
```

where Hour is a logical column, perhaps mapped to a timestamp physical column using the date-and-time Hour(<<timeExpr>>) function.

Rather than entering the numbers 17 and 23 into this expression as constants, you could use the Variable tab of the Variable dialog to set up a static repository variable named `prime_begin` and initialize it to a value of 17, and create another variable named `prime_end` and initialize it to a value of 23.

Static repository variables must have default initializers that are either numeric or character values. In addition, you can use Expression Builder to insert a constant as the default initializer, such as Date, Time, and TimeStamp. You cannot use any other value or expression as the default initializer for a static repository variable.

In previous releases, the Administration Tool did not limit the values of default initializers for static repository variables. Because of this, if your repository has been upgraded from a previous release, you may see warnings in the Consistency Checker similar to the following:

```
The variable, 'Current Month' does not have a constant default initializer.
```

If you see warnings similar to this, update the relevant static repository variables so that the default initializers have constant values.

About Dynamic Repository Variables

You initialize dynamic repository variables in the same way as static variables, but the values are refreshed by data returned from queries. When defining a dynamic repository variable, you create an initialization block or use a preexisting one that contains a SQL query. You also set up a schedule that the Oracle BI Server will follow to execute the query and periodically refresh the value of the variable.

When the value of a dynamic repository variable changes, all cache entries associated with a business model that reference the value of that variable are purged automatically.

Each query can refresh several variables: one variable for each column in the query. You schedule these queries to be executed by the Oracle BI Server.

Dynamic repository variables are useful for defining the content of logical table sources. For example, suppose you have two sources for information about orders. One source contains recent orders and the other source contains historical data.

You need to describe the content of these sources on the Content tab of the Logical Table Source dialog. Without using dynamic repository variables, you would describe the content of the source containing recent data with an expression such as:

```
Orders.OrderDates."Order Date" >= TIMESTAMP '2001-06-02 00:00:00'
```

This content statement becomes invalid as new data is added to the recent source and older data is moved to the historical source. To accurately reflect the new content of the recent source, you would have to modify the fragmentation content description manually. Dynamic repository values can be set up to do it automatically.

Another suggested use for dynamic repository values is in `WHERE` clause filters of logical table sources, defined on the Content tab of the Logical Table Source dialog.

A common use of these variables is to set filters for use in Oracle BI Presentation Services. For example, to filter a column on the value of the dynamic repository variable `CurrentMonth`, set the filter to the variable `CurrentMonth`.

Creating Repository Variables

This section explains how to create repository variables.

To create a repository variable:

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, select **Action > New > Repository > Variable**.
3. In the Variable dialog, type a name for the variable.

Names for all variables should be unique. The names of system session variables are reserved and cannot be used for other types of variables.

4. Select the type of variable: **Static** or **Dynamic**.
5. If you selected **Dynamic**, use the **Initialization Block** list to select an existing initialization block that will be used to refresh the value on a continuing basis.
To create a new initialization block, click **New**. See ["Creating Initialization Blocks"](#) for more information.
6. To add a **Default initializer** value, type the value in the **Default initializer** box, or click the **Expression Builder** button to use Expression Builder.

For static repository variables, the value you specify in the **Default initializer** window persists. It will not change unless you change it. If you initialize a variable using a character string, enclose the string in single quotes (`'`). Static repository variables must have default initializers that are constant values.

7. Click **OK**.

Using Repository Variables in Expression Builder

After they are created, variables are available for use in Expression Builder. In Expression Builder, click the **Repository Variables** folder in the left pane to display all repository variables (both static and dynamic) in the middle pane by name.

To use a repository variable in an expression, select it and double-click. Expression Builder pastes it into the expression at the active cursor insertion point.

Variables should be used as arguments of the function `VALUEOF()`. This happens automatically when you double-click the variables to paste them into the expression.

For example, the following CASE statement is identical to the one explained in the preceding example, except that variables have been substituted for the constants:

```
CASE WHEN "Hour" >= VALUEOF("prime_begin") AND "Hour" < VALUEOF("prime_end") THEN
'Prime Time' WHEN ... ELSE...END
```

Note: You cannot use variables to represent columns or other repository objects.

About Session Variables

Session variables are similar to dynamic repository variables in that they obtain their values from initialization blocks. Unlike dynamic repository variables, however, the initialization of session variables is not scheduled. When a user begins a session, the Oracle BI Server creates new instances of session variables and initializes them.

Unlike a repository variable, there are as many instances of a session variable as there are active sessions on the Oracle BI Server. Each instance of a session variable could be initialized to a different value.

Session variables are primarily used when authenticating users against external sources such as database tables or LDAP servers. If a user is authenticated successfully, session variables can be used to set filters and permissions for that session. For information about using session variables when setting up security, see "Managing Session Variables" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

This section contains the following topics:

- [About System Session Variables](#)
- [About Nonsystem Session Variables](#)

About System Session Variables

System session variables are session variables that the Oracle BI Server and Oracle BI Presentation Services use for specific purposes. System session variables have reserved names that cannot be used for other kinds of variables (such as static or dynamic repository variables and nonsystem session variables).

When you use these variables for Oracle BI Presentation Services, preface their names with `NQ_SESSION`. For example, to filter a column on the value of the variable `LOGLEVEL`, set the filter to the variable `NQ_SESSION.LOGLEVEL`.

[Table 18–1](#) describes the available system session variables.

Table 18–1 System Session Variables

Variable	Description
USER	Holds the value the user enters as his or her logon name. This variable is typically populated from the LDAP profile of the user.
PROXY	Holds the name of the proxy user. A proxy user is a user that has been authorized to act for another user. See <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i> for more information about the PROXY system session variable.

Table 18–1 (Cont.) System Session Variables

Variable	Description
GROUP	Contains the groups to which the user belongs. Exists only for compatibility with previous releases. Legacy groups are mapped to application roles automatically.
WEBGROUPS	Specifies the Catalog groups (Presentation Services groups) to which the user belongs, if any. Note that the recommended practice is to use application roles rather than Catalog groups.
USERGUID	Contains the global unique identifier (GUID) of the user, typically populated from the LDAP profile of the user.
ROLES	Contains the application roles to which the user belongs.
ROLEGUIDS	Contains the global unique identifiers (GUIDs) for the application roles to which the user belongs. GUIDs for application roles are the same as the application role names.
PERMISSIONS	Contains the permissions held by the user, such as oracle.bi.server.impersonateUser or oracle.bi.server.manageRepositories.
DISPLAYNAME	Used for Oracle BI Presentation Services. It contains the name that is displayed to the user in the greeting in the Oracle BI Presentation Services user interface. It is also saved as the author field for catalog objects. This variable is typically populated from the LDAP profile of the user.
PORTALPATH	Used for Oracle BI Presentation Services. It identifies the default dashboard the user sees when logging in (the user can override this preference after logged on).
LOGLEVEL	The value of LOGLEVEL (a number between 0 and 5) determines the logging level that the Oracle BI Server uses for user queries. This system session variable overrides a variable defined in the Users object in the Administration Tool. If the administrator user (defined upon install) has a Logging level defined as 4 and the session variable LOGLEVEL defined in the repository has a value of 0 (zero), the value of 0 applies.
REQUESTKEY	Used for Oracle BI Presentation Services. Any users with the same nonblank request key share the same Oracle BI Presentation Services cache entries. This tells Oracle BI Presentation Services that these users have identical content filters and security in the Oracle BI Server. Sharing Oracle BI Presentation Services cache entries is a way to minimize unnecessary communication with the Oracle BI Server.
SKIN	Determines certain elements of the look and feel of the Oracle BI Presentation Services user interface. The user can alter some elements of the user interface by picking a style when logged on to Oracle BI Presentation Services. The SKIN variable points to an Oracle BI Presentation Services folder that contains the nonalterable elements (for example, figures such as GIF files). Such directories begin with sk_. For example, if a folder were called sk_companyx, the SKIN variable would be set to companyx.
DESCRIPTION	Contains a description of the user, typically populated from the LDAP profile of the user.
USERLOCALE	Contains the locale of the user, typically populated from the LDAP profile of the user.
DISABLE_CACHE_HIT	Used to enable or disable Oracle BI Server result cache hits. This variable has a possible value of 0 or 1.
DISABLE_CACHE_SEED	Used to enable or disable Oracle BI Server result cache seeding. This variable has a possible value of 0 or 1.

Table 18–1 (Cont.) System Session Variables

Variable	Description
DISABLE_SUBREQUEST_CACHE	Used to enable or disable Oracle BI Server subrequest cache hits and seeding. This variable has a possible value of 0 or 1.
SELECT_PHYSICAL	Identifies the query as a SELECT_PHYSICAL query. See "Syntax and Usage Notes for SELECT_PHYSICAL" for more information.
DISABLE_PLAN_CACHE_HIT	Used to enable or disable Oracle BI Server plan cache hits. This variable has a possible value of 0 or 1.
DISABLE_PLAN_CACHE_SEED	Used to enable or disable Oracle BI Server plan cache seeding. This variable has a possible value of 0 or 1.
TIMEZONE	Contains the time zone of the user, typically populated from the LDAP profile of the user.

About Nonsystem Session Variables

You use the same procedure to define nonsystem session variables as for system session variables.

A common use for nonsystem session variables is setting user filters. For example, you could define a nonsystem variable called SalesRegion that would be initialized to the name of the sales region of the user.

You could then set a security filter for all members of a group that would allow them to view only data pertinent to their region.

When you use these variables for Oracle BI Presentation Services, preface their names with NQ_SESSION. For example, to filter a column on the value of the variable SalesRegion, set the filter to the variable NQ_SESSION.SalesRegion.

Creating Session Variables

This section explains how to create session variables.

To create a session variable:

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, select **Action > New > Session > Variable**.
3. In the Session Variable dialog, type a variable name.

Names for all variables should be unique. The names of system session variables are reserved and cannot be used for other types of variables.

4. For session variables, you can select the following options:
 - **Enable any user to set the value.** Select this option to set session variables after the initialization block has populated the value (at user login) by calling the ODBC store procedure NQSSetSessionValue(). For example, this option lets non-administrators to set this variable for sampling.
 - **Security Sensitive.** Select this option to identify the variable as sensitive to security when using a row-level database security strategy, such as a Virtual Private Database (VPD). When filtering cache table matches, the Oracle BI Server looks at the parent database object of each column or table that is referenced in the logical request projection list. If the database object has the **Virtual Private Database** option selected, the Oracle BI Server matches a list of security-sensitive variables to each prospective cache hit. Cache hits would

only occur on cache entries that included and matched all security-sensitive variables.

5. Use the **Initialization Block** list to select an initialization block that will be used to refresh the value on a continuing basis.

To create a new initialization block, click **New**. See ["Creating Initialization Blocks"](#) for more information.

6. To add a **Default Initializer** value, type the value in the **Default Initializer** box, or click the **Expression Builder** button to use Expression Builder.
7. Click **OK**.

Working with Initialization Blocks

Initialization blocks are used to initialize dynamic repository variables, system session variables, and nonsystem session variables. For example, the `NQ_SYSTEM` initialization block is used to refresh system session variables.

This section contains the following topics:

- [About Using Initialization Blocks with Variables](#)
- [Creating Initialization Blocks](#)
- [Associating Variables with Initialization Blocks](#)
- [Establishing Execution Precedence](#)
- [When Execution of Session Variable Initialization Blocks Cannot Be Deferred](#)
- [Enabling and Disabling Initialization Blocks](#)

About Using Initialization Blocks with Variables

An initialization block contains the SQL statement that will be executed to initialize or refresh the variables associated with that block. The SQL statement must reference physical tables that can be accessed using the connection pool specified in the **Connection Pool** field in the Initialization Block dialog.

If you want the query for an initialization block to have database-specific SQL, you can select a database type for that query. If a SQL initialization string for that database type has been defined when the initialization block is instantiated, this string is used. Otherwise, a default initialization SQL string is used.

Caution: By default, when you open the Initialization Block dialog for editing in online mode, the initialization block object is automatically checked out. While the initialization block is checked out, the Oracle BI Server may continue to refresh the value of dynamic variables refreshed by this initialization block, depending on the refresh intervals that are set. When you check in the initialization block, the value of the dynamic variables is reset to the values shown in the Default initializer. If you do not want this to occur, use the **Undo Check Out** option.

This section contains the following topics:

- [Initializing Dynamic Repository Variables](#)
- [Initializing Session Variables](#)

- [About Row-Wise Initialization](#)

Initializing Dynamic Repository Variables

The values of dynamic repository variables are set by queries defined in the **Default initialization string** field of the Initialization Block dialog. You also set up a schedule that the Oracle BI Server will follow to execute the query and periodically refresh the value of the variable. If you stop and restart the Oracle BI Server, the server automatically executes the SQL statements in repository variable initialization blocks, reinitializing the repository variables.

The Oracle BI Server logs all SQL queries issued to retrieve repository variable information in `nquery.log` when the logging level for the administrator account (set upon installation) is set to 2 or higher. You should set the logging level to 2 for the administrator to provide the most useful level of information. You can find the `nquery.log` file in:

```
ORACLE_INSTANCE\diagnostics\logs\OracleBIServerComponent\coreapplication_obisn
```

For more information about user-level logging, see "Managing the Query Log" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Initializing Session Variables

As with dynamic repository variables, session variables obtain their values from initialization blocks. Unlike dynamic repository variables, session variables are not updated at scheduled time intervals. Instead, the Oracle BI Server creates new instances of those variables whenever a user begins a new session. The values remain unchanged for the duration of the session.

Execution of session variable initialization blocks during session logon can be deferred until their associated session variables are actually accessed within the session. See "[Creating Initialization Blocks](#)" for more information.

The Oracle BI Server logs all SQL queries issued to retrieve session variable information if the logging level is set to 2 or higher in the Identity Manager User object, or the `LOGLEVEL` system session variable is set to 2 or higher in the Variable Manager.

The default location for the `nquery.log` file is:

```
ORACLE_INSTANCE\diagnostics\logs\OracleBIServerComponent\coreapplication_obisn
```

For more information about user-level logging, see "Managing the Query Log" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

About Row-Wise Initialization

You can use the row-wise initialization option to create session variables dynamically and set their values when a session begins. The names and values of the session variables reside in an external database that you access through a connection pool. The variables receive their values from the initialization string that you type in the Initialization Block dialog.

For example, suppose you want to create session variables using values contained in a table named `RW_SESSION_VARS`. The table contains three columns:

- `USERID`, containing values that represent the unique identifiers of the users

- NAME, containing values that represent session variable names
- VALUE, containing values that represent session variable values

Table 18–2 shows the table in this example.

Table 18–2 Sample Session Variables Database Table

USERID	NAME	VALUE
JOHN	LEVEL	4
JOHN	STATUS	FULL-TIME
JANE	LEVEL	8
JANE	STATUS	FULL-TIME
JANE	GRADE	AAA

To use row-wise initialization, create an initialization block and select the **Row-wise initialization** option (refer to "Creating Initialization Blocks"). For this example, you would provide the following SQL statement for the initialization string:

```
SELECT NAME, VALUE
FROM RW_SESSION_VARS
WHERE USERID= 'VALUEOF(NQ_SESSION.USERID) '
```

Note that `NQ_SESSION.USERID` has already been initialized using another initialization block.

The following session variables would be created:

- When John connects to the Oracle BI Server, his session contains two session variables from row-wise initialization: `LEVEL`, containing the value 4, and `STATUS`, containing the value `FULL_TIME`.
- When Jane connects to the Oracle BI Server, her session contains three session variables from row-wise initialization: `LEVEL`, containing the value 8; `STATUS`, containing the value `FULL-TIME`; and `GRADE`, containing the value `AAA`.

Initializing a Variable with a List of Values You can also use the row-wise initialization option to initialize a variable with a list of values. You can then use the `SQL IN` operator to test for values in a specified list.

For example, using the table values in the previous example, you would type the following SQL statement for the initialization string:

```
SELECT 'LIST_OF_USERS', USERID
FROM RW_SESSION_VARS
WHERE NAME= 'STATUS' AND VALUE= 'FULL-TIME'
```

This SQL statement populates the variable `LIST_OF_USERS` with a list, separated by colons, of the values `JOHN` and `JANE` (for example, `JOHN: JANE`). You can then use this variable in a filter, as shown in the following `WHERE` clause:

```
WHERE TABLE.USER_NAME = valueof(NQ_SESSION.LIST_OF_USERS)
```

The variable `LIST_OF_USERS` contains a list of values, that is, one or more values. This logical `WHERE` clause expands into a physical `IN` clause, as shown in the following statement:

```
WHERE TABLE.USER_NAME IN ('JOHN', 'JANE')
```

Creating Initialization Blocks

See ["About Using Initialization Blocks with Variables"](#) for more information about initialization blocks.

To create initialization blocks, perform the steps in the following sections:

- [Assigning a Name and Schedule to Initialization Blocks](#)
- [Selecting and Testing the Data Source and Connection Pool](#)

Assigning a Name and Schedule to Initialization Blocks

For repository variables, you can specify the day, date, and time for the start date, as well as a refresh interval.

To assign a name and schedule to initialization blocks:

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, from the **Action** menu, choose **New > Repository** (or **Session**) > **Initialization Block**.
3. In the [Repository | Session] Variable Initialization Block dialog, type a name for the block. (The `NQ_SYSTEM` initialization block name is reserved.)
4. (Repository initialization blocks only) In the **Schedule** area, select a start date and time and the refresh interval.
5. (Session init blocks only) Select the following options when appropriate:

- **Disabled.** If you select this option, the initialization block is disabled.

You can also right-click an existing initialization block in the Variable Manager and choose **Disable** or **Enable**. This option enables you to change this property without opening the initialization block dialog.

- **Allow deferred execution.** If you select this option, execution of the initialization block is deferred until an associated session variable is accessed for the first time during the session.

This option prevents execution of all session variable initialization blocks during the session logon stage, giving a shorter logon time. Session variables that are not needed during the session do not have their initialization blocks executed. This saves the resources which would have been used to execute these unnecessary initialization blocks.

The deferred execution of an initialization block also triggers the execution of all unexecuted predecessor initialization blocks. All associated variables of the initialization block and its unexecuted predecessors are updated with the values returned from the deferred execution.

Note: The **Allow deferred execution** option is unavailable in some circumstances. See ["When Execution of Session Variable Initialization Blocks Cannot Be Deferred"](#) for more information.

- **Required for authentication.** If you select this option, this initialization block must succeed for users to log in. In other words, users are denied access to Oracle Business Intelligence if the initialization block fails to execute. Failure to execute can occur if the wrong credentials have been defined in the initialization block, or if there is an error in the default initialization string.

Note that this requirement is waived for internal processes (like Delivers) that use impersonation, if a single user session variable has been associated with

the initialization block. In this case, the trusted internal process can connect regardless of whether the initialization block succeeds or fails.

The next step is to select the data source and connection pool.

Selecting and Testing the Data Source and Connection Pool

If you select **Database** as the data source type for an initialization block, the values returned by the database for the columns in your SQL statement are assigned to variables that you associate with the initialization block. For session variable initialization blocks, you can also select **LDAP Server** or **Custom Authenticator**.

It is recommended that you create a dedicated connection pool for initialization blocks where you select **Database** as the data source type. See "[About Connection Pools for Initialization Blocks](#)" for more information.

If you select **Database** as the data source type:

- If you select **Database** as the data source type, and do not select the **Use OBI EE Server** option

The SQL statement used to refresh the variable must reference physical tables that can be accessed through the connection pool specified in the **Connection Pool** field. The tables do not have to be included in the Physical layer of the metadata. At run time, if an initialization string for the database type has been defined, this string is used. Otherwise, the default initialization SQL for the database type is used. You can overwrite this string.

When you create SQL and submit it directly to the database (for example, when using database-specific SQL in initialization blocks), the SQL statement bypasses the Oracle BI Server. The order of the columns in the SQL statement and the order of the variables associated with the initialization block determine which columns are assigned to each variable.

You should test this SQL using the **Test** button in the [Repository | Session] Variable Initialization Block Data Source dialog. If the SQL statement contains an error, the database returns an error message. See "[Testing Initialization Blocks](#)" for more information.

- If you select **Database** as the data source type, and select the **Use OBI EE Server** option

The SQL statement you use to refresh the variable might be written for a specific database. However, it will still work with other data sources because the SQL statement is processed by the Oracle BI Server. The Oracle BI Server can also provide functions (such as `PI`) that might not be available in the data source, and the SQL statement will work with other data sources supported by the Oracle BI Server (for example, ADF, SQL Server, Oracle, and XML files). When you select the **Use OBI EE Server** option, there is no need for a connection pool, because the SQL statement is sent to the Oracle BI Server and not directly to the underlying database.

You can only test this SQL statement using the **Test** button in the [Repository | Session] Variable Initialization Block Data Source dialog when in online mode. If the SQL statement contains an error, the database returns an error message. See "[Testing Initialization Blocks](#)" for more information.

To select a data source and connection pool for initialization blocks:

1. In the Administration Tool, select **Manage**, then select **Variables**.

2. In the Variable Manager dialog, double-click the initialization block you want to edit. You can edit Repository initialization blocks, or Session initialization blocks.
3. Click **Edit Data Source** next to the **Connection Pool** field.
4. From the **Data Source Type** list, select one of the following types.
 - **Database:** For repository and session variables.
 - **LDAP Server:** For session variables.
 - **Custom Authenticator:** For session variables. See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for more information.
5. If you selected **Database** for your data source type, perform one of the following steps:
 - Select **Default initialization string** or **Use database specific SQL**, and then perform the following steps:
 - a. Click **Browse** next to the **Connection Pool** field to select the connection pool associated with the database where the target information is located. If you do not select a connection pool before typing the initialization string, you receive a message prompting you to select the connection pool.
 - b. In the Select Connection Pool dialog, select the connection pool and click **Select**. You must select a connection pool before typing an initialization string.

By default, the first connection pool under the database object in the Physical layer is not available for selection. This behavior ensures that you cannot use the same connection pool for initialization blocks that you use for queries. See "[About Connection Pools for Initialization Blocks](#)" for more information

You can change this behavior so that the first connection pool is available for selection by selecting **Allow first Connection Pool for Init Blocks** in the Options dialog, although this is not recommended. See "[Setting Preferences](#)" for more information.
 - c. If you selected **Use database specific SQL**, then in the **Database** pane, expand and select the database. Then, enter its associated string.

Otherwise, in the **Default initialization string** box, type the SQL initialization string needed to populate the variables. See "[Examples of Initialization Strings](#)" for examples.
 - d. (Optional) Click **Test** to test the data source connectivity for the SQL statement.
 - e. Click **OK** to return to the Initialization Block dialog.
 - Select **Use OBI EE Server**, and then perform the following steps:
 - a. In the box, enter the SQL initialization string needed to populate the variables.

The string you enter here is processed by the Oracle BI Server, and therefore as long as it is supported by the Oracle BI Server, the string will work with different data sources.

For example, an initialization block might use the function `pi()`, which is specific to SQL Server. However, if you select **Use OBI EE Server**, the query is rewritten by the Oracle BI Server for the appropriate database. In

other words, if you change the SQL Server back-end database to Oracle, the query will still work.

See ["Examples of Initialization Strings"](#) for additional examples.

- b. Click **OK** to return to the Initialization Block dialog.
6. If you selected **LDAP Server** for your data source type, perform the following steps:
 - a. Click **Browse** to select an existing LDAP Server, or click **New** to open the General tab of the LDAP Server dialog and create an LDAP Server.
 - b. Click **OK** to return to the Initialization Block dialog.

The LDAP server name and the associated domain identifier appear in the **Name** and **Domain identifier** columns.
7. If you selected **Custom Authenticator** for your data source type, perform the following steps:
 - a. Click **Browse** to select an existing custom authenticator, or click **New** to create one.
 - b. Click **OK** to return to the Initialization Block dialog.
8. Click **OK**.

Examples of Initialization Strings This section contains the following initialization string examples:

- [Example 18–1, "A SQL Statement When Site Uses Delivers"](#)
- [Example 18–2, "A SQL Statement When Site Does Not Use Delivers"](#)
- [Example 18–3, "A SQL Statement Joining Tables From Multiple Data Sources - When Using the 'OBI EE Server' Setting"](#)

Example 18–1 A SQL Statement When Site Uses Delivers

```
SELECT username, groupname, dbname, schemaname FROM users
WHERE username=:USER'
NQS_PASSWORD_CLAUSE (and pwd=:PASSWORD')NQS_PASSWORD_CLAUSE
```

This SQL contains two constraints in the WHERE clause:

' :USER ' (note the colon and single quotes) is the ID the user types when logging in.

' :PASSWORD ' (note the colon and single quotes) is the password the user enters. This is another system variable whose presence is always assumed when the USER system session variable is used. You do not need to set up the PASSWORD variable, and you can use this variable in a database connection pool to allow passthrough login using the user ID and password of the user. You can also use this variable in a SQL statement.

When using external table authentication with Delivers, the portion of the SQL statement that makes up the :PASSWORD constraint must be embedded between NQS_PASSWORD_CLAUSE clauses.

The query returns data only if the user ID and password match values found in the specified table. You should test the SQL statement outside of the Oracle BI Server, substituting valid values for the USER and PASSWORD variables and removing the NQS_PASSWORD_CLAUSE clause.

For more information, see *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Example 18–2 A SQL Statement When Site Does Not Use Delivers

```
SELECT username, groupname, dbname, schemaname FROM users
WHERE username=:USER'
AND pwd=:PASSWORD'
```

This SQL statement contains two constraints in the WHERE clause:

' :USER ' (note the colon and the single quotes) is the ID the user types when logging in.

' :PASSWORD ' (note the colon and the single quotes) is the password the user enters. This is another system variable whose presence is always assumed when the USER system session variable is used. You do not need to set up the PASSWORD variable, and you can use this variable in a database connection pool to allow passthrough login using the user ID and password of the user. You can also use this variable in a SQL statement.

The query returns data only if the user ID and password match values found in the specified table. You should test the SQL statement outside of the Oracle BI Server, substituting valid values for the USER and PASSWORD variables.

Example 18–3 A SQL Statement Joining Tables From Multiple Data Sources - When Using the 'OBI EE Server' Setting

```
select WUSER.name, wuser_detail.email
from "db-11g/orcl"."NAME"."WUSER',
"sqlxpress"."master"."dbo"."wuser_detail"
where username=:USER:
```

The above query example in the initialization block uses a join query with multiple tables from different data sources (for example, SQLServer, Oracle and XML Files). The query works because when you select the **Use OBI EE Server** option, the query is rewritten by the BI Server for the specified data sources.

Testing Initialization Blocks You should test the SQL statement using the **Test** button or a SQL tool such as the Oracle BI Client utility. If you use a SQL tool, be sure to use the same DSN or one set up identically to the DSN in the specified connection pool.

In online mode, Initialization Block tests do not work with connection pools set to use :USER and :PASSWORD as the user name and password. In offline mode, the Set values for variables dialog is displayed so that you can populate :USER and :PASSWORD.

To test initialization blocks (optional):

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, double-click the initialization block.
3. In the [Repository | Session] Variable Initialization Block dialog, click **Edit Data Source**.
4. In the [Repository | Session] Variable Initialization Block Data Source dialog, click **Test**.

Note: The **Test** button is disabled when the **Use OBI EE Server** option is selected in offline mode.

5. In the Set value for the variables dialog, verify the information is correct, and then click **OK**.
6. In the View Data from Table dialog, type the number of rows and the starting row for your query, and then click **Query**.

The Results dialog lists the variables and their values.

The next step is to associate variables with the initialization block.

Associating Variables with Initialization Blocks

The SQL `SELECT` statement in the Default initializer list can contain multiple columns. The order of the columns in the SQL statement and order of the variables associated with the initialization block determine the column value that is assigned to each variable. Therefore, when you associate variables with an initialization block, the value returned in the first column is assigned to the first variable in the list.

For repository variable initialization blocks, when you open a repository in online mode, the value shown in the **Default initialization string** field of the Initialization Block dialog is the current value of that variable as known to the Oracle BI Server. The number of associated variables can be different from the number of columns being retrieved. If there are fewer variables than columns, extra column values are ignored. If there are more variables than columns, the additional variables are not refreshed (they retain their original values, whatever they may be). Any legal SQL can be executed using an initialization block, including SQL that writes to the database or alters database structures, assuming the database permits the user ID associated with the connection pool to perform these actions.

If you stop and restart the Oracle BI Server, the server automatically executes the SQL statement in the repository variable initialization blocks, re-initializing the repository variables.

For session variable initialization blocks, you can select Row-wise initialization. The **Use caching** option is automatically selected when you select the **Row-wise initialization** option. Selecting the **Use caching** option directs the Oracle BI Server to store the results of the query in a main memory cache. See "[About Row-Wise Initialization](#)" for more information.

The Oracle BI Server uses the cached results for subsequent sessions. This can reduce session startup time. However, the cached results might not contain the most current session variable values. If every new session needs the most current set of session variables and their corresponding values, you should clear this option.

See "[About Using Initialization Blocks with Variables](#)" for more information.

To associate variables with initialization blocks:

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, double-click the initialization block you want to edit. You can edit repository initialization blocks, or session initialization blocks.
3. Click **Edit Data Target**.
4. In the [Repository | Session] Variable Initialization Block Variable Target dialog, perform one of the following steps:
 - Associate variables with the initialization block by doing one of the following:
 - Click **New**, and in the Variable dialog, create a new variable. See "[Creating Repository Variables](#)" or "[Creating Session Variables](#)" for information about creating variables.

- Click **Link** to associate an existing variable with an initialization block. Then, in the **Browse** dialog, select the variable to be refreshed by this initialization block and click **OK**.

Note: For the **Custom Authenticator** data source type (Session variables only), the variable **USER** is required.

- Select **Row-wise initialization**. This option is for session variable initialization blocks only. See "[About Row-Wise Initialization](#)" for more information. If you select **Row-wise initialization**, the **Use caching** option becomes available.
5. To reorder variables, select a variable and click **Up** or **Down**.
 6. To remove a variable from association with this block, select the variable and click **Remove**.
 7. Click **OK**.

The next step is to establish execution precedence.

Establishing Execution Precedence

When a repository has multiple initialization blocks, you can set the order (establish the precedence) in which the blocks will be initialized.

First, you open the block that you want to be executed last and then add the initialization blocks that you want to be executed before the block you have open. For example, suppose a repository has two initialization blocks, A and B. You open initialization block B, and then specify that block A will execute before block B. This causes block A to execute according to block B's schedule, in addition to its own.

To establish execution precedence:

1. In the Administration Tool, select **Manage**, then select **Variables**.
2. In the Variable Manager dialog, double-click the last initialization block that you want to be initialized.
3. In the [Repository | Session] Variable Initialization Block dialog, click **Edit Execution Precedence**.
4. In the [Repository | Session] Variable Initialization Block Execution Precedence dialog, click **Add**.
Add is only available if there are initialization blocks that have not yet been selected.
5. In the Browse dialog, select the blocks that should be initialized before the block that you have open, and then click **OK**.
6. To remove a block, in the [Repository | Session] Variable Initialization Block Execution Precedence dialog, select the block you want to remove and click **Remove**.
7. Click **OK**.
8. If you want the initialization block to be required, in the [Repository | Session] Variable Initialization Block dialog, select the **Required for authentication** option.
9. Click **OK**.

Note: When you select the **Use OBI EE Server** option for an initialization block:

- Execution precedence does not apply, because during user login, an initialization block with the **Use OBI EE Server** option selected is executed after initialization blocks with the **Use OBI EE Server** option not selected.
 - The **Required for authentication** option is dimmed, because this type of initialization block is executed after authentication.
-
-

When Execution of Session Variable Initialization Blocks Cannot Be Deferred

Execution of session variable initialization blocks cannot be deferred in some circumstances, including when row-wise initialization is being used and when the **Required for authentication** option has been selected. When the execution of session variable initialization blocks cannot be deferred, a message is displayed that explains why. See "[Assigning a Name and Schedule to Initialization Blocks](#)" for more information.

The following list summarizes the scenarios in which execution of session variable initialization blocks cannot be deferred:

- The **Row-wise initialization** option is selected in the Session Variable Initialization Block Variable Target dialog
 Example message: "The execution of init block 'A_blk' cannot be deferred as it is using row-wise initialization."
- The **Required for authentication** option is selected in the Session Variable Initialization Block dialog.
 Example message: "The execution of init block 'A_blk' cannot be deferred as it is required for authentication."
- The **Data Source Type** is not Database.
 Example message: "The execution of init block 'A_blk' cannot be deferred as it does not have a connection pool."
- The initialization block is used by session variables named `PROXY` or `USER`.
 Example message: "The execution of init block 'A_blk' cannot be deferred as it is used by session variable 'PROXY'."
- The initialization block is used by session variables where the **Security Sensitive** option is selected in the Session Variable dialog.
 Example message: "The execution of init block 'A_blk' cannot be deferred as it is used by session variable 'A' which is security sensitive."
- The initialization block is a predecessor to another initialization block which does not have the **Allow deferred execution** option selected.
 Example message: "One of the successors for init block 'A_blk' does not have "Allow deferred execution" flag set. Init block 'B_blk' does not have "Allowed deferred execution" flag set."

Enabling and Disabling Initialization Blocks

You can use the Variable Manager in the Administration Tool to enable and disable initialization blocks.

To enable or disable an initialization block:

1. In the Administration Tool, select **Manage**, then select **Variables**. The Variable Manager appears.
2. In the left pane, select **Initialization Blocks** under **Repository** or **Session**, depending on whether you want to enable or disable repository initialization blocks or session initialization blocks.
3. In the right pane, right-click the initialization block you want to enable or disable.
4. Choose **Enable** or **Disable** from the right-click menu.
5. Close the Variable Manager and save the repository.

Managing the Repository Lifecycle in a Multiuser Development Environment

This appendix provides best practice information for managing the lifecycle of the Oracle BI repository when you are using a multiuser development environment.

Building your Oracle BI repository using the multiuser development environment enables you to do the following:

- Build large, interrelated semantic models
- Independently build multiple, independent semantic models to run in the same Oracle BI Server and Presentation Services server
- Develop several branches on different schedules, in parallel, while fixing urgent bugs or enhancement requests on the production version
- Incrementally design and test at the individual and team levels
- Enable individual developers to design and test manageable subsets without impacting each other, yet share their changes with other developers in a controlled, incremental fashion
- Migrate changes to test and production systems in bulk, or incrementally

This appendix covers the development lifecycle of the Oracle BI repository (RPD). It does not cover the development lifecycle of the Oracle BI Presentation Catalog used by Presentation Services. This appendix also does not cover how to use the multiuser development environment for Independent Software Vendor (ISV) organizations building portable BI applications for sale as products.

See also [Appendix B, "MUD Case Study: Eden Corporation"](#) for detailed examples of how the multiuser development environment is used in a typical business scenario.

This appendix contains the following topics:

- [Planning Your Multiuser Development Deployment](#)
- [Multiuser Development Architecture](#)
- [Understanding the Multiuser Development Environment](#)
- [MUD Tips and Best Practices](#)
- [Troubleshooting Multiuser Development](#)

Planning Your Multiuser Development Deployment

This section describes tasks you need to perform as part of the planning phase before beginning multiuser development.

This section contains the following topics:

- [About Business Organization and Governance Process Best Practices](#)
- [About Technical Team Roles and Responsibilities](#)

About Business Organization and Governance Process Best Practices

You need to provide a strong, effective governance process to make decisions about shared resources and to resolve conflicts among the many stake-holders. As in any business process, you must have a strong business sponsor, and the steering committee must be staffed with strong business people who can negotiate effectively and make good decisions that will not change over time. Having an effective governance process has proven to be the single most important factor in achieving successful multiuser development with Oracle Business Intelligence.

Before you begin your multiuser development project, you must first lay out the business value, priorities, roadmap, and requirements, as well as lower level details of the design, as described in [Table A-1](#).

Table A-1 Tasks to Accomplish During the Planning Phase

For...	You must...
Strategic requirements	<ul style="list-style-type: none"> ■ Determine which business processes to measure ■ Determine which data sources and subject areas to access
Business requirements for repository objects	<ul style="list-style-type: none"> ■ Select and define metrics, dimensions, and hierarchies ■ Identify objects that will be shared between development teams ■ Resolve conflicts between teams ■ Define Presentation layer subject areas
Security requirements	<ul style="list-style-type: none"> ■ Define Application Roles and corresponding privileges for your user base ■ Define which repository developers can access which metadata and data
Development	<ul style="list-style-type: none"> ■ Determine the styles of multiuser development to use ■ Define areas to break down into MUD projects ■ Determine the owners for metadata objects
Project management	<ul style="list-style-type: none"> ■ Set initiatives - purpose, goals, requirements, scope, schedule, budget ■ Define phases - scope, schedule ■ Allocate resources - hardware, software, databases, developers ■ Decide on a strategy for development branching ■ Prioritize and schedule production updates from different development teams
Operations	<ul style="list-style-type: none"> ■ Negotiate service level agreements ■ Coordinate schedules for updates and downtime

About Technical Team Roles and Responsibilities

This section describes the hands-on roles involved in repository development and its lifecycle. Depending on the size of your company and team, some of these roles might be served by one person.

Repository development roles include:

- MUD administrator (one for each development team, plus backup)
 - Assigns repository password
 - Sets up and maintains MUD projects
 - Manages the master repository shared directories
 - Manages branches and branch merges
 - Manages repository migrations
 - Manages test and production connection pools
 - Manages independent semantic models (has metadata read/write privileges for all models)
- Repository developer (many per development team)
 - Knows the repository password
 - Owns, operates, and maintains a personal development sandbox that includes all necessary Oracle Business Intelligence components
 - Manages user and application role provisioning on their sandbox stack
 - Creates functional and data authorization content in the repository
 - Performs unit testing
 - Performs check-outs, and check-in merges and publishing, as required
- Production Operations staff
 - Knows the repository password (for managing connection pools and applying patches)
 - Applies updated repositories, and applies XML patch updates to the running BI Server's repository
 - Can log in to production computers and read/write the Oracle Business Intelligence directories or run programs
 - Manages the production file system, including the repository directory, logs, and configuration files
 - Manages the production servers (Administration Server, Managed Servers with Java components, and Oracle Business Intelligence system components like Oracle BI Server and Presentation Services)
 - Manages production security, including provisioning users, groups, and application roles
 - Manages and migrates application roles in production
 - Manages production connection pools (in the case where the MUD administrator does not have security privileges for production connection information)

People in other roles outside the repository development team are also involved. These include people administering the test environment and running the tests, and also the Oracle BI Presentation Catalog developers.

Multiuser Development Architecture

Before you can understand best practices for MUD and repository lifecycle management, you need to understand the architecture of the development environment.

This section contains the following topics:

- [About Multiuser Development Concepts](#)
- [About Multiuser Development Styles](#)
- [Multiuser Development Sandbox Architecture](#)
- [Multiuser Development and Lifecycle Management Architecture](#)

About Multiuser Development Concepts

This section explains fundamental concepts related to developing and deploying systems for multiuser development.

Repository (RPD)

The repository, or RPD file, is the fundamental artifact under development. It defines all the metadata used by the Oracle BI Server for interpreting user requests, applying role-based security, generating queries to data sources, and post-processing the results.

Application Roles and the Policy Store

A secondary artifact under development is the set of application roles. User object permissions, data access filters, and query limits (governors) are defined against these application roles in the repository logic. Presentation Services also uses application roles for assigning its privileges and permissions.

You can use the default policy store embedded in Oracle WebLogic Server, or you can use a separate external policy store. If you are using the embedded policy store, you define application roles in Fusion Middleware Control, which persists them in the Policy Store in Oracle WebLogic Server. You can then use the Administration Tool in online mode to add application roles from your policy store to your repository at design time. At run time, the Oracle BI Server uses the application roles provisioned to each user to apply the correct security privileges to user requests.

Sandboxes, Projects, and Branches

An instance of the repository is usually edited by only one repository developer at a time. Multiple developers work in parallel on subset instances of the repository, called "projects." They work in separate sandbox environments, and merge their changes into a master repository instance frequently to distribute changes and pick up changes made by others in the team. This approach enables the creation of very large enterprise applications. It also enables independent semantic models to be developed by separate teams and merged into the master repository for production hosting in a single Oracle BI Server cluster. Finally, it enables branching and merging so that teams can work on major projects in parallel, and can even make emergency fixes to the main code line in production without disrupting ongoing development projects.

You typically use the Simple install type when installing a development sandbox.

Single, Shared Repository

Presentation Services connects to just one repository that has been uploaded to the Oracle BI Server. The metadata for all semantic models must reside in this single

repository, even if the semantic models share no objects. See "[About Multiuser Development Styles](#)" for more information about semantic models in a repository.

Repository Password

The repository file is protected by the repository password. The Oracle BI Server needs this password to open and load the repository at startup. It stores the repository password in the secure Credential Store. You must also enter this password when you open the repository in the Administration Tool or other utilities and line commands. Note that user logon credentials are stored in the identity store, not the credential store.

Oracle BI Presentation Catalog

The Oracle BI Presentation Catalog is an important BI application artifact that contains the metadata that defines reports, dashboards, KPIs, scorecards, and other reporting layer objects. The Oracle BI Presentation Catalog is outside the scope of this document. See "Managing Objects in the Oracle BI Presentation Catalog" in *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition* for more information about Oracle BI Presentation Catalog development.

Migration

The completed repository is migrated to test and production systems using Fusion Middleware Control. No downtime is necessary, because you can refresh clustered production Oracle BI Servers with a rolling restart.

Deployment Parameters During Migration

Some repository parameters must change when migrating a repository between development, test, and production systems, such as connection pool settings. These parameters must change because they are based on the deployment, not the application logic. You can automate these updates using the Oracle BI Server XML API (`biserverxmlexec.exe -B`). During multiuser development, developers merging in content are automatically prevented from overwriting the master repository test connection pool and database parameters with their local unit test parameters.

Application Role (Policy Store) Migration

There are several options for migrating application roles between development, test, and production systems. For simplicity, this document assumes you will re-key a small number of application role names by hand. For full information about migrating application roles, and other migration considerations, see "Moving Oracle Business Intelligence Components to a Production System" in *Oracle Fusion Middleware Administrator's Guide*.

Users and the Identity Store

As a best practice, users are not represented by metadata objects in the repository at design time. Also, the repository does not manage or store their credentials. Instead, users must always be provisioned to application roles in the run-time environment to receive privileges. Their credentials, as well as their mapping to application roles through groups, are managed in an external Identity Store. See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Setting the Environment for the Administration Tool and Utilities

The current directory structure generalizes the location of metadata, data, and configuration files. For this reason, each program or utility you launch requires an

initial setup of the configuration for the Oracle instance to which that program belongs. For command-line utilities, you must first launch the bi-init utility to set the parameters, and then launch the utility from a command line inside the bi-init window. Most Oracle BI Server tools, including the Administration Tool, run the bi-init utility transparently. If you run a command without setting the instance environment first, the program will typically fail to find external files.

About Multiuser Development Styles

Choose your style of development based on the size of your team, the number of teams and parallel initiatives, and your requirements for security and availability. [Table A-2](#) shows the multiuser development styles.

Table A-2 *Styles of Multiuser Development*

Style	Description
Serial Development (Figure A-1)	You can use this method if you have a small number of developers and low concurrency. Development users share a repository file through e-mail, a shared directory, or on a shared development system, and only one of them makes changes at a time. They must coordinate with each other on the development schedule.
Serial Development with Patch Files (Figure A-1)	As a variation on serial development, you can share a base binary repository, and ship changes only between users using patch files.
Shared Online Development (Figure A-2)	The best practice is for only one developer at a time to develop metadata in online mode against a single Oracle BI Server and its repository. However, multiple online users are an option for development situations where communication among the team members is frequent, a higher risk of conflicts is acceptable, and minimum administrative overhead is a goal.
MUD (Figure A-3)	The Multiuser Development feature enables over one hundred development users to work in parallel on a shared, enterprise repository. Each user can develop and unit test in a separate sandbox environment, using only manageable-sized subsets of the metadata. When a unit of work is complete, they can automatically merge it into the branch, where other users can pick up those changes and integrate them with their own metadata. When a project phase is ready for promotion, the MUD administrator migrates it to the test environment, and eventually, production. The MUD administrator manages branches and sub-branches to enable parallel development of independent initiatives or fixes, and merges them into the main branch to incrementally migrate them to test and production environments. The MUD administrator also manages fine-grained "projects," which are the manageable-sized repository subsets individual developers check out to their local sandbox environments. See " Understanding the Multiuser Development Environment " for additional information.

Table A-2 (Cont.) Styles of Multiuser Development

Style	Description
MUD with Multiple, Independent Semantic Models (Figure A-3 and Figure A-4)	There might be cases where you need two or more independent semantic models, rather than a single, integrated, enterprise-wide model. This is common due to security requirements, or when unrelated divisions of a business share a common Oracle Business Intelligence infrastructure. The MUD administrator creates a branch for each model, which enables parallel development and integrated testing for each team's semantic model. When an independent semantic model's branch is ready for promotion to production, the MUD administrator simply merges the branch into main. The MUD administrator can set security on the branches so that each developer can only see the semantic model to which they are assigned, and so that only the MUD administrator and selected production operations staff can access the integrated main model.
MUD with Delegated Administration (Figure A-3 and Figure A-4)	When the independent semantic models are developed by different organizations on different schedules, a centralized MUD administrator might not provide the desired level of local control. In this case, you can provide a dedicated MUD administrator for each independent semantic model's branch. The branch administrator operates in the same way as an ordinary MUD administrator. In this scenario, the MUD super-administrator defines a branch for each organization, checks out the subset repository, and provides it to the branch administrator. When the model is ready for promotion to production, the branch administrator passes the repository back to the super-administrator, who merges it into the main branch for promotion, and then migrates the combined repository to production.

Figure A-1 the serial development style of multiuser development.


Figure A-1 Serial Development

Figure A-2 shows the shared online development style of multiuser development.

Figure A-2 Shared Online Development




Figure A-3 shows true multiuser development with branching.

Figure A-3 Multiuser Development




Figure A-4 shows the architecture for a repository with multiple, independent semantic models.

Figure A-4 Repository with Multiple, Independent Semantic Models




Table A-3 shows which multiuser development styles meet various requirements for security and availability.

Table A-3 Requirements Met by Multiuser Development Styles

Requirement	Serial	Shared Online	MUD with Single Semantic Model	MUD with Multiple Semantic Models	MUD with Delegated Administration
No administrator	Yes	No	No	No	No
Up to five concurrent developers	No	Yes	Yes	Yes	Yes
More than five concurrent developers	No	No	Yes	Yes	Yes
Work on manageable subsets of a large repository, such as Oracle BI Applications	No	No	Yes	Yes	Yes
Built-in checkout, merge, and rollback	No	No	Yes	Yes	Yes
Host independent semantic models in single repository	Yes	Yes	No	Yes	Yes
Incremental migration of units of work to production	No	No	Yes	Yes	Yes
Developers of independent semantic models cannot see each others' metadata	No	No	No	Yes ¹	Yes ¹
Each independent semantic model has its own MUD administrator	No	No	No	No	Yes

¹ Requires secure MUD Directory. An overall MUD administrator must still have access to all metadata from all teams.

Multiuser Development Sandbox Architecture

When using MUD, each developer works on their own, fully dedicated sandbox Oracle Business Intelligence system. You should set up your sandbox to contain all the components you need for development and unit testing.

First, you need to decide whether to use a UNIX or Windows server for your Oracle Business Intelligence stack. Follow these guidelines:

- If you choose the Windows-only option, make sure your system has enough memory. Note that you will need additional resources if you choose to host your database on the same hardware. See "[System Requirements and Certification](#)" for information about minimum hardware requirements.
- If you choose the UNIX option, you still need a Windows system to run the Administration Tool. Use the Oracle Business Intelligence Simple install type on the UNIX system, and use the Client install type on the Windows system to install the Administration Tool.

Note: The Client install type does not include Catalog Manager, so if you plan to develop the Oracle BI Presentation Catalog in this environment, you should use the Simple install type on the Windows computer rather than the Client install type.

In online mode, the Oracle BI Server loads the repository from its local repository directory on the UNIX system in:

```
ORACLE_INSTANCE/bifoundation/OracleBIServerComponent/coreapplication_  
obisn/repository
```

Note that the Administration Tool on Windows also points to a local /repository directory by default, but you can use any directory for offline development.

You also need to install a development database. This database can be a dedicated, personal database, or it can be shared among multiple repository developers. Note the following considerations about the development database

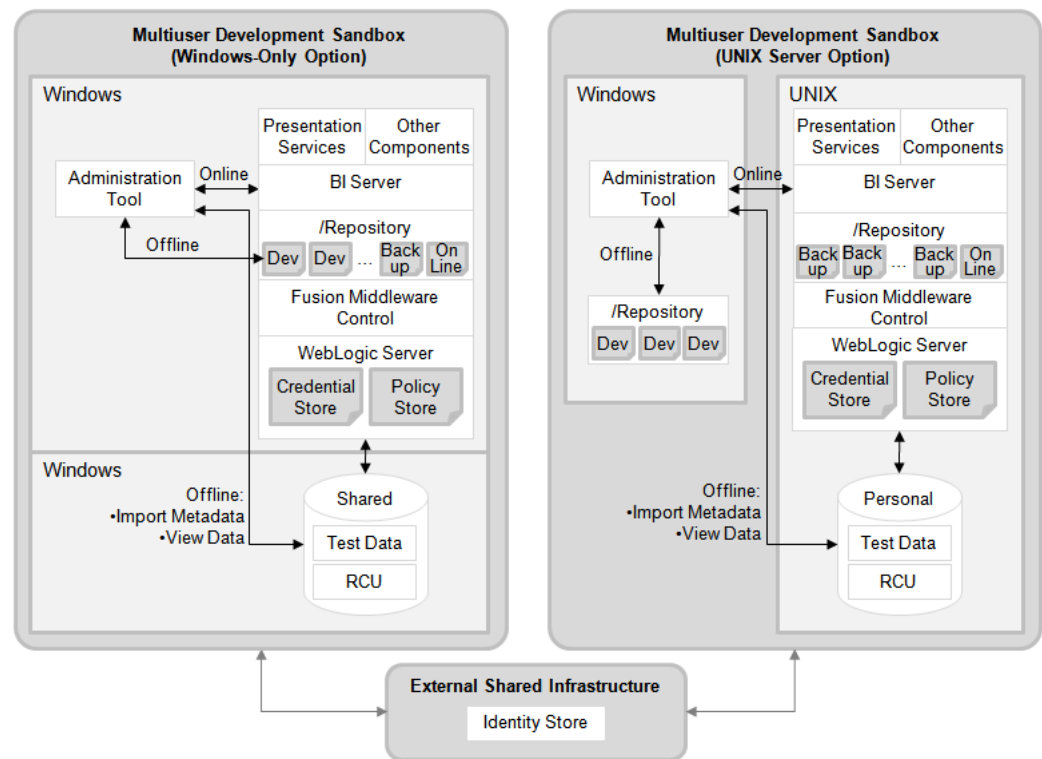
- **Platform:** You can choose to host your development database on your sandbox computer if you provide enough memory, or you can host it on a centralized, shared server. Both scenarios are shown in [Figure A-5](#).
- **RCU:** The database must contain the schemas required by Oracle Business Intelligence. You load these schemas using the Repository Creation Utility (RCU). These schemas enable support for Oracle BI Scheduler and annotations for Oracle Scorecard and Strategy Management, provide sample tables for Usage Tracking, and enable many other features. The Oracle WebLogic Server Managed Servers for Oracle Business Intelligence, and all the services that depend on it, require access to a running database with the required RCU schemas in order to operate.
- **Data Source Schemas:** You also need data source schemas for the metadata under development. You can optionally include some data source schemas in your RCU database, or they can be in other databases. Note the following additional information about data source schemas:
 - **Test Data:** The data source schemas should be loaded with test data. If users are testing read-only metadata, the schemas can be shared among multiple development sandboxes. They can be located on the development sandbox computer if enough memory is available.
 - **Multiple Sources:** Optionally, your environment might include multiple data sources needed by your initiative, such as other relational sources, Essbase, Hyperion Financial Management, Microsoft Analysis Services, SAP B/W, and others. These sources can be shared or dedicated, local or remote.

- **Connectivity:** You must set up connectivity from your Administration Tool and Oracle Business Intelligence stack to each data source. This configuration can include installing the required drivers or clients, setting up ODBC DSNs, setting up native connectivity, and other steps. See [Chapter 4, "Importing Metadata and Working with Data Sources"](#) and [Chapter 15, "Setting Up Data Sources on Linux and UNIX"](#) for full information.

Note that for Oracle Database connectivity, Oracle Business Intelligence requires an instance of `TNSnames.ora` in `ORACLE_HOME/network/admin`.

Figure A-5 shows the architecture of the multiuser development sandbox.

Figure A-5 Multiuser Development Sandbox Architecture




Note: Most developers prefer to disable caching in the development sandbox. This makes it easier to validate and debug physical queries using the log. When the cache is enabled, the physical SQL might not appear in the log, because the request might get fulfilled by the cache. In this release, you must disable caching using Fusion Middleware Control. See "Using Fusion Middleware Control to Enable and Disable Query Caching" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Multiuser Development and Lifecycle Management Architecture

The overall MUD architecture contains the developer sandbox systems described in the previous section, as well as any test and production systems. In addition, there are several additional major components, as follows:

- **Windows MUD administration system.** This system is maintained by the MUD administrator. Note the following about this system:
 - It provides one shared network MUD directory for the main branch, and additional shared network MUD directories for each side branch. The Windows permissions on each shared directory only allow access to the developers for that branch. Each shared directory stores the master repository for that branch, as well as various control and history files for MUD functions.
 - It has a client installation of Oracle Business Intelligence. The Administration Tool and Oracle BI Server utilities are used for creating and managing MUD projects, performing merges, creating patches, and other MUD administrator tasks. Other Oracle Business Intelligence processes like the Oracle BI Server, as well as the policy store and credential store, are typically not used on this platform.
 - A 32 bit or 64 bit system can be used, because none of the Oracle Business Intelligence Java components, system components, or other infrastructure are used on this computer.
- **One or more test systems.** These systems are used for running integrated tests of merged content. They run the full Oracle Business Intelligence stack, and can be either UNIX- or Windows-based. These systems are frequently clustered.
- **Oracle BI Presentation Catalog system.** Optionally, you might have a system with a full Oracle Business Intelligence stack for developing Oracle BI Presentation Catalog content.
- **Clustered production system.** Eventually, you will have a clustered production system on one of the supported Oracle Business Intelligence platforms.
- **External identity store.** This appendix assumes you are using an external identity store like Oracle Internet Directory.

Figure A-6 shows a sample deployment architecture for the repository lifecycle using the multiuser development environment.

Figure A-6 Example Multiuser Development Deployment Architecture

Understanding the Multiuser Development Environment

MUD is a set of features that enables teams of developers to work in parallel on projects of any size, despite the complex interrelationships and dependencies in the repository model. With MUD, you can:

- Divide the repository file into subsets
 - Enables users to work with manageable subsets when the repository is very large
 - Enables independent testing for each subset by each developer or team
 - Makes it easier to manage merges later after checking out a branch subset
 - Enables you to separate independent semantic models into secure branches for development
- Incrementally develop, test, and migrate
- Merge subsets and branches, handling conflicts between user changes
- Apply Oracle updates to a packaged BI Application you have modified
- Merge separately developed applications into a single repository
- Access history logging and audit information
- Roll back to historical repository states

The multiuser development feature also provides the following other useful capabilities:

- Coordinates merging into the master, including tracking original repository files
- Provides locking for reliable updates
- Logs changes
- Automatically backs up repositories before each potentially destructive operation

This section contains the following topics:

- [About Multiuser Development Environment Task Flow](#)
- [About Multiuser Development Projects](#)
- [How to Create Branches](#)
- [Which Merge Utility Should I Use?](#)

About Multiuser Development Environment Task Flow

The basic flow of working with multiple users is as follows:

1. A developer defines the "starter" Physical layer, as well as basic facts and subject areas. This provides some basic objects to anchor the MUD projects.
2. The MUD administrator defines projects and puts the RPD into the main branch MUD directory. Note that the MUD directory where the master repository is stored cannot be the same as the Oracle BI Server local repository directory.
3. A developer can now check out one or more projects, do development work, and then check back in by merging the changes into the master, and then publishing it back to the MUD directory.
4. Meanwhile, other developers check out and do development on the same or other projects (note that projects are for subsetting, not for enforcing security). Because check-in uses a three-way merge, users can check out, develop, and check in their changes in any order. Even property changes to a single object from multiple users are merged. If conflicts do occur between users, the three-way merge feature provides a way for the developer to choose which objects to keep. Communication between users is a key to avoiding and resolving conflicts, and you should have your governance process assign ownership of major objects in order to avoid such conflicts.
5. When a development phase is complete, the MUD administrator can migrate the content to a test system. There might be several iterations back through check out, bug fix, check in, and retest. When the repository passes the testing phase, the MUD administrator can migrate it to the production environment.
6. The MUD administrator can create and manage multiple development branches as large MUD projects. A branch can be secured to ensure that only one development team can work on it. A branch can even be treated recursively as a main, with its own, delegated MUD administrator.

About Multiuser Development Projects

The multiuser development feature is built around a metadata object called a *project*. The project is the unit of check-out from the master repository, and the subsequent check-in merge and publish. When a master repository becomes very large, a project is a manageable-sized subset that a developer can check out to work on. It is also designed to be self-consistent, so that you can run the consistency checker (analogous

to compile-time code checking) and then test it on the Oracle BI Server with a client such as Answers at run time. When you are satisfied with the results, you can merge it back into the master repository so that it becomes part of the larger application. Meanwhile, history is logged and repository backups are automatically created at key points.

MUD features in the Administration Tool streamline this flow for fine-grained developer projects. Similarly, superset projects streamline the management and merging of branches.

The project subset contains a set of metadata objects. You define a project to include a minimum set of objects explicitly, but many others are included implicitly. Having objects implicitly added to projects simplifies your project management task.

The following objects are explicitly specified by the MUD administrator as members of a project:

- Logical fact tables
- Presentation layer subject areas
- Application roles
- Users (although the best practice is to only use application roles in RPD logic)
- Initialization blocks
- Variables

All other objects are implicitly included in a project and are found by the Administration Tool during the check-out process. For example:

- Descendents of the explicitly defined objects. For example, when a logical fact table is included explicitly, all its logical columns are included in the project implicitly.
- Logical dimension tables that join to the selected logical fact tables, and the join objects themselves.
- Logical table sources for the included logical fact and dimension tables.
- Physical tables that map to the logical tables included in the project, and the joins between them.
- Marketing target levels and list catalogs.

Note that objects that are in the list of explicitly defined objects are sometimes included implicitly. For example, if a logical column contains an expression that includes a variable, the variable is implicitly included in the project, even if the MUD administrator does not explicitly add it.

It is normal for projects to overlap. An object that appears in one project can also appear in another project. For example, some customers prefer to create an overall project for each independent semantic model, as well as smaller projects within each independent model for checking out individual units of development work. You can also check out multiple projects simultaneously to work on a larger set of metadata.

[placeholder for diagram of how a project draws in some dimensions and physical objects]

See also "[Setting Up Projects](#)" for additional information.

How to Create Branches

This section explains how to create main branches, side branches, and delegated administration branches.

This section contains the following topics:

- [How to Create a Main Branch](#)
- [How to Create a Side Branch](#)
- [How to Create a Delegated Administration Branch](#)

How to Create a Main Branch

The ultimate master repository is usually source-controlled in the main branch, out of which all branches and ultimately all development projects check out. The main branch usually stages the repository in production. That is, to migrate content to production, you merge it into the main branch, and then migrate the main repository to the production system.

Similarly, to fix a production bug, a developer typically checks out of the main branch. The developer then fixes the bug, and then merges it back into the main branch for migration to test and production. Meanwhile, parallel development in side branches is not affected.

To create the main branch as the MUD administrator, you must first create a shared directory and copy the master repository file to it. The directory can be on either Windows or UNIX, but the UNIX share must be accessible by Windows users.

Set the security on the share to only allow access by the appropriate developers. Depending on your requirements, you might only allow developers to access the side branch master repositories, not the main branch master repository.

If this is a new project, you typically have a developer seed the repository with initial content that can be easily split into branch projects.

How to Create a Side Branch

The best practice for branching is to start with a superset MUD project, and then use the MUD check-out, merge, and publish features. Then, individual users or sub-branches use finer-grained projects and check out of the branch master. Using MUD for this functionality provides automatic back-ups at the check-out points, tracks original repositories to ensure correct merges, uses more optimistic merge assumptions that require less user intervention, and provides history and roll-backs.

To create a side branch as the MUD administrator:

1. In the main master repository, create a project that extracts all content required for the branch. Follow these guidelines to create the project:
 - If little or no metadata has been designed in the repository, it is a best practice to seed it with content that can anchor the project. This makes it easier to ensure the project extracts the physical content you need to support the logical fact tables. Usually, this means one or more logical fact tables are created, with at least some representative columns. The columns should be mapped to the physical tables and joins needed to support the fact tables. Finally, create the project and define the objects that belong to it.
 - If content already exists, create the project and define the objects needed in that branch. The branch can overlap with other projects, if necessary.

- It is also possible to create an empty project for check-out. However, the developer who checks it out must ensure that all the physical objects that need to be implicitly added to the project are mapped to the logical fact table before check-in. Similarly, the developer must ensure dimensions are joined before check-in to ensure their inclusion, and must explicitly add any subject areas, variables, initialization blocks, application roles, and users. This method is more prone to errors than seeding the project before defining it.
- Typically, connection pools for environments such as production must be secured. Ensure that the connection pool settings in the master repository are acceptable for the developers to access. Note that developers typically change the settings to match their local test databases. At check-in, connection pool and database settings are not merged, to prevent overwriting the settings in the master repository.

Use the Oracle BI Server XML API to automate connection pool changes required during migrations to production and other environments. See "Moving from Test to Production Environments" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

2. Create a shared MUD directory for the branch. Every branch should have its own MUD directory. Set the permissions so that only the developers working on that branch have access to it.

Note that you can use branch permissions combined with project subsetting to prevent developers from seeing metadata that belongs to other teams. Design the projects carefully so that they only extract metadata related to one team. This goal is easiest to achieve if the teams use different business models, subject areas, physical models, variables, initialization blocks, and application roles.

It is also a best practice to use a consistent system of naming and numbering your branches.

3. Check out the branch project using **File > Multiuser > Checkout**. You can check out into your local repository directory, or another directory.
4. Copy the repository to the branch MUD directory, where it serves as the master repository.
5. Define fine-grained MUD projects for developers to check out from the branch. Inform the developers that the branch is ready for development.
6. Based on your project plan, your developers perform a final check-in (merge and publish) of their changes when they have completed development and unit tests.
7. When all check-ins of planned content are complete for the phase, the branch project is ready to undergo integrated testing. To accomplish this, migrate the branch master repository file to the test environment. When a bug is found, the assigned developer checks out the appropriate projects, fixes the bug, and tests the metadata. After the changes are checked in, migrate the branch repository to the test environment again. Note that this branch project can be tested without impacting, or being impacted by, development work in other branches.
8. When integrated testing is complete, the branch is ready to promote to production. Remove the branch master repository from the branch shared directory so that users cannot change it. Copy it back into your local repository directory, and merge it into the main using the Administration Tool. The main repository is now ready for migration to integrated test and production.

9. Typically, the MUD administrator checks out the branch again and places the branch repository in the shared MUD directory for the next phase of development. Note that during the check-out, any changes from other branches, or bug fixes from the main branch, are picked up by the branch repository.

How to Create a Delegated Administration Branch

You can use a branch to delegate local control of a metadata subset to the organization that is developing and maintaining it. To do this, you assign a branch MUD administrator to the branch, who performs the same roles as the main MUD administrator. This approach works best with an independent semantic model, so that you can ensure that there is no metadata overlapping with other groups.

The delegated branch MUD administrator performs the same tasks as the main branch administrator, including defining projects for further branches and creating fine-grained projects for developers.

To create a delegated administration branch as the main MUD administrator:

1. Set permissions on the main MUD directory so that only the main MUD administrator (and the main MUD administrator backup) have access.
2. Create a branch MUD project, branch MUD directory, and checked-out branch master repository as described in the previous section.
3. Set security on the branch MUD directory so that the main MUD administrator and the delegated branch MUD administrator have access.
4. The branch administrator defines projects for further branches, as well as fine-grained projects for developers. If required, the branch administrator deploys additional branches off the delegated branch for development initiatives, with permissions set to allow developers to check out of these repositories.
5. Developers fix production bugs by checking out of the delegated branch MUD directories, because individual developers are not allowed access to the main branch.
6. When developers check in all their changes, the branch administrator checks their branches into the delegated branch for integrated testing.
7. To promote a delegated branch to production after integrated testing is complete, the main MUD administrator performs the following two steps:
 - a. Removes the branch master repository from the delegated branch repository shared directory and checks it back into the main branch using the Administration Tool.
 - b. Migrates the main branch master repository to production.
8. Typically, the main MUD administrator checks out the branch again and places the branch repository in the delegated branch shared MUD directory for the next phase of development. The branch administrator then checks out next-level branches and places their repositories into the branch shared MUD directories, so that developers can check out their fine-grained projects and begin their work.

Which Merge Utility Should I Use?

There are several different merge tools that are optimized for various situations and environments. When deciding which merge approach and utility to use, you should consider whether you need to perform the task on Windows or UNIX systems. You should also consider your other requirements, such as whether you need to merge

changes you made to a semantic model, or whether you need to combine two semantic models from different development efforts.

[Table A-4](#) shows which merge approaches and tools meet various requirements.

Table A-4 Requirements Met by Different Merge Approaches

Requirement	Merge Approach	Tools Used	Platform
<ul style="list-style-type: none"> ■ Merge a checked-out MUD project back into master repository ■ Merge a checked-out MUD branch project back into the main branch master repository 	Three-way merge	<ul style="list-style-type: none"> ■ MUD merge 	Windows
<ul style="list-style-type: none"> ■ Combine non-MUD branches and changes back into the main branch 	Three-way merge	<ul style="list-style-type: none"> ■ Merge Repository Wizard (Full Merge selected) 	Windows
<ul style="list-style-type: none"> ■ Apply an Oracle update XML patch to customized, deployed BI Application ■ Apply an update XML patch you created from development to a deployed repository 	Three-way merge	<ol style="list-style-type: none"> 1. Merge Repository Wizard (Patch Merge selected) 2. Patchrpd utility 	<ol style="list-style-type: none"> 1. Windows 2. All
<ul style="list-style-type: none"> ■ Combine disjoint logical content with potential ID conflicts 	Two-way merge	<ul style="list-style-type: none"> ■ Merge Repository Wizard (with blank original) 	Windows
<ul style="list-style-type: none"> ■ Combine disjoint content guaranteed in advance by the developer to have no conflicts (all platforms) 	Insert-Update-Delete	<ul style="list-style-type: none"> ■ biserverxmlxec -B ■ biserverxmlcli (online) ■ Copy/Paste XML 	All
<ul style="list-style-type: none"> ■ Combine disjoint content guaranteed in advance by the developer to have no conflicts (Windows only) 	Insert-Update-Delete	<ul style="list-style-type: none"> ■ Copy/Paste Administration Tool tool objects ■ Administration Tool Import from Repository (deprecated) 	Windows

See ["Merging Repositories"](#) for more information about merging.

MUD Tips and Best Practices

This section provides tips and best practices for working in a multiuser development environment.

This section contains the following topics:

- [Best Practices for Branching](#)
- [Best Practices for Setting Up Projects](#)
- [Best Practices for Three-Way Merges](#)
- [Best Practices for MUD Merges](#)
- [Best Practices for Two-Way Merges](#)
- [Best Practices for Production Migration](#)
- [Best Practices for Application Roles and Users](#)

Best Practices for Branching

Follow these guidelines for creating side branches:

- The MUD directory where the master repository is stored cannot be the same as the Oracle BI Server local repository directory.
- A branch should be a checked-out MUD project. This automates and streamlines many of the tasks of merging the branch back into the main branch, such as using the correct original repository.
- Always put the checked-out branch master repository into its own MUD directory. Then, let developers check out their fine-grained projects from the branch master repository. When branch development, check-ins, and testing are complete, remove the master from the branch repository directory and check it back into the main branch master repository using the Administration Tool. Then, check it out again and place the new version in the branch MUD directory for development of the next phase.
- Use Windows permissions on the branch MUD directory to control which developers have access to it.
- Set multiuser development options by creating an .opt file in the branch MUD directory. As a best practice, define specific administrators, and set Mandatory Consistency Check and Equalize During Merge to Yes. See "[Setting Multiuser Development Options](#)" for more information.
- Plan your branches based on the increments of functionality you want to deliver to production. Each branch should contain an increment you want to migrate as a unit.
- If you accidentally merge branches in the wrong order, you can roll them back using the MUD history. See "[Viewing and Deleting History for Multiuser Development](#)" for more information.

Best Practices for Setting Up Projects

Follow these guidelines for setting up projects:

- Break your RPD down into fine-grained projects, as small as possible while still being useful. Doing so improves performance and ease of management.
- Break your logical fact tables down into smaller partitions to enable smaller, separate projects.
- For each side branch, overlay a larger project that will extract the branch's contents. This enables the project to manage the checkout and merge of the branch, including tracking of the original repository. Individual developers can check out their development projects from the checked-out branch project. Be sure that all development projects are checked back into the side branch before merging it back into the main branch.
- When you add new content to a repository, be sure it is part of your project before you check it in. If you check in objects that are not part of a project, they will not be in your extract the next time you check the project out. You or the MUD Administrator must then edit the entire repository, or at least several other projects that do happen to include your new content, and then add the objects to the project at that time.
- Sometimes, you might need to extract several projects at the same time to get all the content you need.

Tip: Presentation layer objects, including subject areas, presentation tables, and presentation hierarchies, are now objects that you explicitly include in the project. Unlike in previous releases, the security settings of the Administration Tool user have no impact on which subject areas, presentation tables, or presentation columns are included in a project when checking it out. Instead, the set of Presentation layer objects determines the scope of the project.

See also ["Setting Up Projects"](#) for more information.

Best Practices for Three-Way Merges

Follow these guidelines when performing three-way merges:

- Ensure that you have the original repository from which both the modified and current repositories were built.

Note that this step is done for you automatically in a MUD merge.

- Typically, you should open the development repository as current, then use the main repository as modified, and the starting point of the branch as original.
- Unit test before merging.
- As a best practice, select **Equalize during merge** and **Check consistency of the merged RPD** in the Merge Repository Wizard. See ["Equalizing Objects"](#) for full information about the importance of equalizing objects.

Best Practices for MUD Merges

Follow these guidelines when performing MUD merges:

- Unit test before merging.
- Unit test after merging, but before publishing. Keep in mind that you are holding the lock on the master repository, so keep it brief.
- Be sure your full name is correct in the Tools > Options > Multiuser tab. Doing so assists in logging and in checking who holds the locks.
- When performing the local merge, be sure to write useful comments in the Lock Information screen. You or other administrators can use the comments later to help identify historical repositories when you need to perform rollbacks or other tasks.
- When the MUD administrator is editing the master RPD, it must be inaccessible to checkout users. To accomplish this, you can temporarily remove it from the shared directory and place it in another directory, or you can rename it before editing. Make sure to restore it when the edits are complete.

You can also open the repository in offline mode so that other users are locked out by the Windows file system. Note that you should only use this method when you are sure you will finish all your work in one atomic session.

- Merge frequently. The list of conflicts and decisions needed in a small merge is easy to understand. When the merge is too large, the number of changes make it much harder to understand, and it is much harder to avoid human errors. If you need to roll back, the number of changes discarded is also much bigger. Performance is also better for small merges.

- If performance of merges is a problem, consider breaking the project down into several, finer-grained projects. Also, be sure to merge more frequently, so the number of changes in the merge is smaller and therefore faster.
- Because local connection pool changes are overridden by the connection pool settings in the master repository at each check-in time, the local test settings must be reapplied at each checkout if they are different from the master. It is best to automate application of your local connection pool settings using the Oracle BI Server XML API. See "Moving from Test to Production Environments" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.
- The most successful large teams have formal process requirements and expectations for the communications and tasks between the repository developers and the MUD administrator. For example, they might have a form for a request for the MUD administrator to create a project. Many teams also have service level agreements and lead times, such as 24 hours to create a project.
- Set the option to force a consistency check during MUD merges. A clean consistency check ensures that the Oracle BI Server can load the model correctly, similar to the way a compiler checks to see if code can be generated correctly. Even if the merge seems to succeed, an inconsistent repository may have trouble with starting the Oracle BI Server, online repository editing check-ins, and subsequent merges. See "[Setting Multiuser Development Options](#)" for information about how to enable this feature.
- Set the option to force an equalize before merge. This reduces the number of duplicate objects, since it is common for developers to import the same physical tables, for example. See "[Setting Multiuser Development Options](#)" for information about how to enable this feature.

See also "[Equalizing Objects](#)" for full information about the importance of equalizing objects.
- Do not delete or change content needed by others, unless you are the owner and have coordinated with the other developers. If you delete a column you do not need in your project, that action usually causes it to be deleted from the master when you merge, even if other users depend on it.

Tip: Presentation object aliases receive special treatment in merges. Their purpose is to hold historical names of objects, so that when names change, old reports do not break. If you changed any names during development, new aliases were added. During merge, you have the option whether to keep any new aliases you have created, or not. You also have the option to keep any or all past aliases, because the historical reports might still exist.

See also "[About the Multiuser Development Merge Process](#)" for more information.

Best Practices for Two-Way Merges

Use two-way merge when you need to combine two repositories that were developed separately into a single repository. This situation usually occurs when you need to host two semantic models in a single repository.

Follow these guidelines when performing two-way merges:

- Make sure that the top-level objects in each repository have different names, so there are no unintentional renames or object merges. Check the following objects:
 - Business models
 - Subject areas
 - Physical databases
 - Variables
 - Initialization block
 - Application roles
 - Users
 - Marketing objects
- Equalize before merging. Doing so honors the fully qualified names over which you have control, and assigns upgrade IDs to ensure there will be no conflicts between the two repositories. See also "[Equalizing Objects](#)" for full information about the importance of equalizing objects.
- In the Administration Tool, perform a full merge with a blank repository as the original file.

To create a blank repository, open a new repository, and save it without importing a source or creating any objects. Although this repository contains default security objects, these do not impact your merges.

Caution: Do not use features like Import from Repository or copy/paste in the Administration Tool to move metadata incrementally. These approaches do not correctly merge changes.

Using these features might produce the results you expect most of the time, but this is just good luck. The rest of the time, values of the upgrade IDs in the metadata objects will clash, effectively causing overwrites of random objects. However, you might not notice the problem until much later when you do regression testing. Because upgrade IDs are assigned sequentially, and are only unique within one repository, clashes are very likely.

You should also use caution when using the `biserverxmlcli` and `biserverxmlexec -b` utilities. Be sure to fully understand the information about managing IDs described in "About Using the Oracle BI Server XML API to Merge and Append Objects" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition*.

See also "[Performing Full Repository Merges Without a Common Parent](#)" for more information about two-way merges.

Best Practices for Production Migration

Follow these guidelines when moving from test to production:

- When updating metadata on the production cluster, perform a rolling restart to restart one Oracle BI Server at a time, so that users do not experience down time while changes are being loaded. You can use the BI Systems Management API to

programmatically start and stop Oracle BI Servers, or you can restart each Oracle BI Server manually in Fusion Middleware Control.

For more information, see "Starting and Stopping Oracle Business Intelligence" and "Starting and Stopping Oracle Business Intelligence Using the Oracle BI Systems Management API" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

- It is not recommended to alter metadata in online mode in production using the Administration Tool.
- It is not recommended to update metadata in online mode in production using the biserverxmlcli utility.

Best Practices for Application Roles and Users

Follow these guidelines when working with application roles and users:

- Do not build data access security around user objects in the repository. Instead, define repository permissions, filters and governors based on application roles.
- The set of application roles should be defined by the governance committee. The business team knows what the business roles are, who is allowed to see which data, and which roles are the same from one application to the next. Therefore, the governance committee is in a position to define and name the application roles and decide which roles can be shared across applications.
- When you create a new Application Role, be sure to add it to a project so that you can check it out again after you merge. Also, if you create a placeholder application role in the Administration Tool in offline mode, make sure to add it to the policy store later.
- You can find whether the application roles used by a repository are provisioned in the system by opening your repository in the Administration Tool in online mode and running the consistency checker. It is recommended that you perform this check each time you migrate the repository and application roles to a new system.
- If you only need to migrate a small number of application roles between environments, you can enter them manually in Fusion Middleware Control on the target system if you are using the embedded policy store in Oracle WebLogic Server.

Troubleshooting Multiuser Development

This section describes common problems and how to resolve them.

Orphan Lock Held on Master RPD

If a user sets a lock by issuing the command to merge local changes, it is not cleared until the user publishes or discards their changes. If the user forgets and leaves for a two-week vacation, the MUD administrator can release the lock.

The lock is stored in a hidden system file in the master directory. If you cannot see the lock file, in Windows Explorer, select **Tools**, then select **Folder Options**. In the View menu, ensure that the option **Show hidden files and folders** is selected.

The lock file has the same name of the master RPD with a .lck extension. Delete the lock file to release the lock on the repository.

[Figure A-7](#) shows a repository lock file.

Figure A-7 Repository Lock File

Name	Size	Type
sales.mhl	5 KB	MHL File
sales.lock	1 KB	LCK File
modified subset of sales.007	28 KB	007 File
sales.rpd	38 KB	Oracle BI Repositor...
sales.006	38 KB	006 File
sales.006.csv	1 KB	Microsoft Office Ex...
modified subset of sales.006	31 KB	006 File

Object Deleted By Other User

If another MUD developer deletes an object that you need, you can choose one of the following options:

- Roll back to an earlier version, and reapply all the changes since then. The easiest way to roll back is generally to replay history in the history log. To do this, choose **File > Multiuser > History**, and then select an entry and use **Actions > View**.

See "[Viewing and Deleting History for Multiuser Development](#)" for more information.

- Re-create the deleted objects, and equalize so that future merges treat it as the same object.

Project Missing Needed Physical Tables and Joins After Checkout

Physical objects do not explicitly belong to a project. Instead, the physical objects mapped to the logical fact tables in your project are extracted at the time of check out.

To get needed physical objects into your local extract, check out an additional project that does have mappings to the physical objects you need. If there is no such project, then the entire repository must be edited to create mappings to a logical fact table in your project. The MUD administrator can take the repository off line to make that change. Then, your next check out should include the physical objects.

Objects Added in the Last Session Missing from Checked Out Repository

If recently added objects are missing from your checked out repository, you might have forgotten to add the objects to your project before you merged and published. Only objects in your project, or inferred from your project (like dimensions and physical objects), are included in your extracted repository.

To resolve this issue, ask the MUD administrator to add the objects to your project in the master repository, and then check out again.

Object Renamed by Appending #1

This situation occurs when two objects are merged with the same fully qualified name, but with different internal upgrade IDs. The merge logic in this situation determines that the objects are semantically different objects, and changes the name of the object to preserve uniqueness.

To resolve this issue, run the equalizerpds utility, which reassigns upgrade IDs so that objects with the same fully qualified names in the two repositories have the same Upgrade IDs. Then, try the merge again. The two objects should be merged instead of causing a rename.

See "[Equalizing Objects](#)" for more information.

Rolling Back to Previous Versions

The multiuser development environment stores backup copies of RPDs at appropriate checkpoints. Each time a potentially destructive operation is performed, a new backup is stored in the master directory. It has the name of the RPD, and the extension is a three-digit incrementing integer. Individual developers can also make copies of their RPD files at any time during development.

In the developer's sandbox, the original version of a checked-out project is stored with the name `originalrpd_name.rpd`. This version is automatically used if the developer discards changes.

You can also view and roll back to an older version by following these steps:

1. Open the Administration Tool, but not a repository.
2. Select **File > Multiuser > History**.
3. Select the version of interest, and then choose **Actions > View > Repository**.
4. Select **File**, then select **Copy As** to save that version to a new name.
5. Use the older version to replace the latest version, or replace the master repository with the older version.

Example

This example explains how to copy an older version to replace the latest version. Assume you are at version 1000 and want to roll back to version 900. In this situation, you have three files: `repository.900`, `repository.1000`, and `repository.rpd`, the current version. To perform the roll back, make a copy of `repository.900` and rename it to `repository.1001`. (This lets you keep `repository.1000` in your version history.) Then, copy `repository.900` to `repository.rpd`.

MUD Case Study: Eden Corporation

This appendix describes a fictional case study that shows how the multiuser development environment might be used for a particular business case.

This appendix contains the following topics:

- [About the Eden Corporation Fictional Case Study](#)
- [Phase I - Initiating Multiuser Development \(MUD\)](#)
- [Phase II - Branching, Fixing, and Patching](#)
- [Phase III - Independent Semantic Model Development](#)

About the Eden Corporation Fictional Case Study

Eden Corporation (a fictional company) recently purchased Oracle Business Intelligence. They have two divisions that are licensed and plan to use the product. Because of this, the company has two separate initiatives:

- **Initiative S:** The Sales Division wants to use Oracle Business Intelligence for dashboarding and analysis of revenue versus plan. They want to deploy an initial phase to production quickly to meet an immediate need. Then, they want to roll out more functionality in Phases II and III. Initiative S is large enough that they will have two developers working on it.
- **Initiative H:** The Human Resources Division (HR) needs to do dashboarding and analysis of HR data. Initiative H is a smaller initiative, so it will have only one developer. They plan to deliver their application to production between Initiative S Phases II and III.

Note that the Sales developers and the HR developers are not allowed to see each others' data or metadata. The metadata administrator is the only person who has security privileges for all the metadata.

As in all organizations, there will also be a steady stream of urgent requests and occasional bugs from production. The developers will need to deliver fixes for these within days, even though the longer-term initiatives S and H are in development at the same time.

About the Technical Team Roles and Responsibilities

Eden Corporation has staffed the team as follows:

- Adam Straight - MUD Administrator
- Sally Andre - Developer for Sales Division, Revenue project
- Scott Baker - Developer for Sales Division, Quota project

- Helen Rowe - Developer for HR Division

About the Eden Corporation Development Phases

Eden Corporation plans to deploy RPDs to production based on the following timeline:

1. January - Sales Phase I (projects Revenue and Quota)
2. February - Sales Phase II (add project Target, extend projects Revenue and Quota)
3. March - HR (one project used)
4. April - Sales Phase III (extend all three projects)

About the Eden Corporation Topology

Eden Corporation plans to use the following systems for their multiuser development environment:

- MUD Administrator - NT computer with a share
- Sally Andre - NT computer for Administration Tool client, and Linux computer to run the Oracle Business Intelligence stack
- Scott Baker - high-powered NT computer
- Helen Rowe - either of the above
- Test - Linux computer
- Production - Clustered Linux computers

About the Repository Architecture

Because of Eden Corporation's business structure and initiatives, they need to have two independent semantic models in their repository: one for Sales and one for HR. Each of these models can have multiple projects.

Planning the Repository Structure

Eden Corporation knows that it is important to plan the structure of their repository file so that it will be able to support the multiuser development needs of their organization. They assigned owners to major objects, so the developers know who to go to when conflicts arise, and which objects they should not modify on their own.

Tip: When hosting multiple independent semantic models, be sure to itemize the names of top-level objects to prevent duplicate names.

Table B-1 and Table B-2 show the high-level repository objects in main.rpd for both Initiative S and Initiative H, mapped to projects and owners. Note that Adam is the overall owner of both Initiative S and Initiative H.

Table B-1 Initiative S Repository Objects Mapped to Projects and Owners

Object Type	Object	Owner	ProjRevenue	ProjQuota	ProjTarget
physical database	Sample App Data	Sally	Yes	Yes	Yes
business model	Sales	Sally	n/a	n/a	n/a
logical fact table 1	F10 Billed Rev	Sally	Yes	Yes	No
logical fact table 2	F30 Facts Targets	Scott	No	No	Yes
logical fact table 3	F50 Facts Quotas	Scott	No	Yes	No
logical dimension	(various)	Sally	Yes	Yes	Yes

Table B-1 (Cont.) Initiative S Repository Objects Mapped to Projects and Owners

Object Type	Object	Owner	ProjRevenue	ProjQuota	ProjTarget
subject area (1)	Sales Quota	Scott	No	Yes	No
subject area (2)	Sales Revenue	Sally	Yes	No	No
subject area (3)	Sales Target	Scott	No	No	Yes
variable	S_Last_Load	Sally	Yes	Yes	Yes
initialization block	S_Last_Load	Sally	Yes	Yes	Yes
application role (1)	Sales Management	Sally	Yes	Yes	Yes
application role (2)	Sales Rep	Sally	Yes	Yes	Yes

Table B-2 Initiative H Repository Objects Mapped to Projects and Owners

Object Type	Object	Owner	ProjHR
physical database	Human Resources Data	Helen	Yes
business model	HR	Helen	n/a
logical fact table (1)	Payroll Facts	Helen	Yes
logical fact table (2)	Medical Ins Facts	Helen	Yes
logical dimension	(various)	Helen	Yes
subject area (1)	HR Payroll	Helen	Yes
subject area (2)	HR Medical	Helen	Yes
variable	H_Last_Load	Helen	Yes
initialization block	H_Last_Load	Helen	Yes
application role (1)	HR Management	Helen	Yes
application role (2)	HR Rep	Helen	Yes

Phase I - Initiating Multiuser Development (MUD)

In the first phase, both Sally Andre and Scott Baker will develop in parallel. Sally will create the starter content, which Adam Straight will divide into projects. He will then create the MUD directory so that Sally and Scott can check out and perform their development. After unit testing, they merge and publish their changes, and then Adam migrates the repository to the test environment. After a bug fix cycle, Adam promotes the repository to production.

The following sections describe Phase I development:

- [Starting Initiative S](#)
- [Setting Up MUD Projects](#)
- [First Developer Checks Out](#)
- [Second Developer Checks Out](#)
- [First Developer Checks In](#)
- [Second Developer Checks In](#)
- [MUD Administrator Test Migration Activities](#)
- [Phase I Testing](#)

- [Phase I Migration to Production](#)
- [Phase I Summary](#)

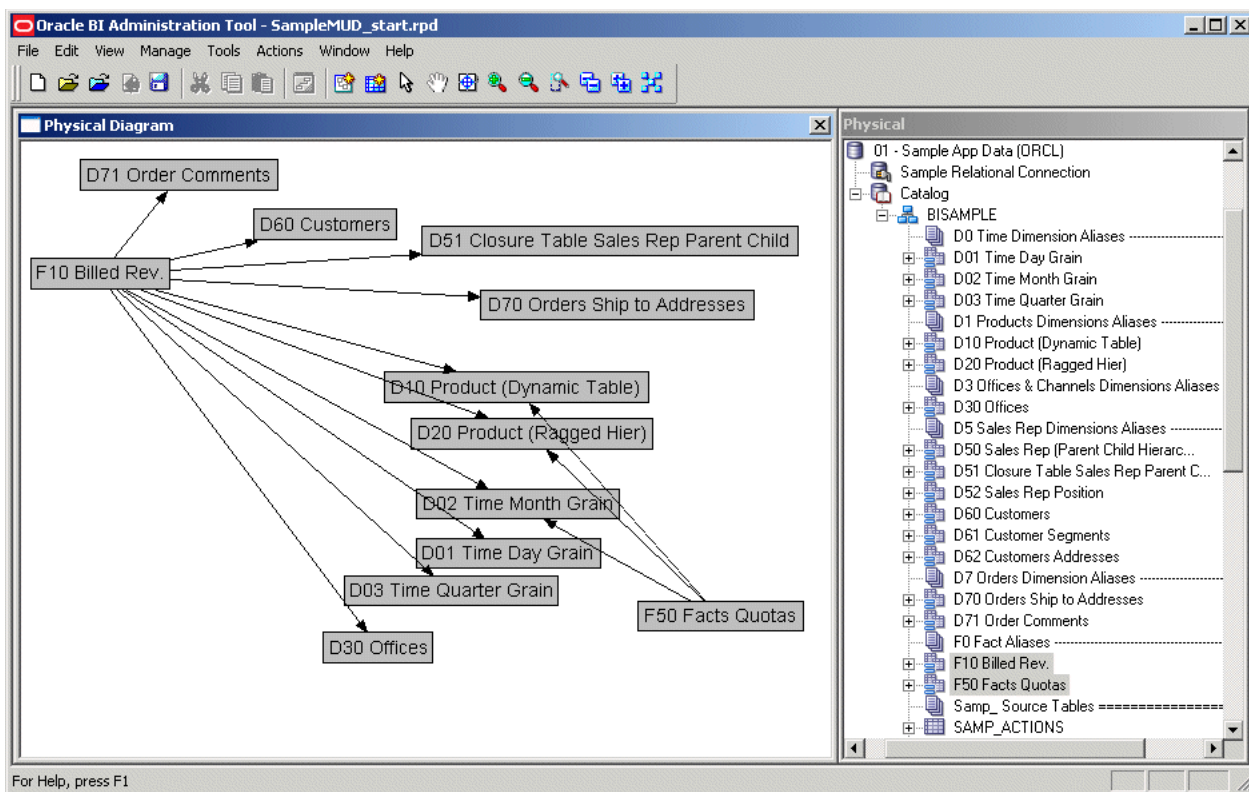
Starting Initiative S

Sally Andre starts off Initiative S from an empty RPD. Because it is easier to divide the repository into MUD projects if you define some logical stars and subject areas first, she begins by developing the physical model needed for Phase I. She includes connection pool details for her own local test data sources.

Tip: The physical model should include the physical tables, the best practice of aliasing all the physical tables to give them meaningful names, and joins.

Figure B-1 shows the physical model for Initiative S.

Figure B-1 Initiative S Physical Model



Sally drags the Physical layer to the Business Model and Mapping layer to create some starter content. She removes unneeded tables, and ensures that the star joins are correct. She also ensures that all the physical tables that will be needed during development have mappings from the starter logical tables, so that they will be included in the projects when they are checked out. For Sally, these steps create two logical fact tables, F10 Revenue and F50 Quotas, that can act as the basis for the projects.

Sally also needs to have some subject areas to map to the projects in the business model. She could drag the entire business model, but a convenient way to accomplish this is to instead right-click the business model and select **Create Subject Areas for Logical Stars and Snowflakes**. This feature creates a subject area from each logical fact table.

Sally does not need to be concerned about the contents of the subject areas yet. All that matters is that each subject area maps to the logical fact table for the same project. However, she does name the subject areas based on the plan agreed to in the governance meeting: Sales Quota and Sales Revenue.

Sally now has enough content for the MUD administrator to create the first two projects based on the Revenue and Quota fact tables. To review, Sally has made sure that she meets the following criteria at a minimum:

1. At least one logical fact table according to the governance plan, to anchor the projects. The columns of the logical fact tables need not be complete or even properly named, but they do need to be complete enough to map all the physical content.
2. Enough logical dimensions so that the repository will pass the consistency check.
3. Physical content that maps to one or more logical fact tables, so they will be included in projects.
4. The subject areas needed according to the governance plan.

Setting Up MUD Projects


The MUD administrator for Eden Corporation, Adam Straight, now handles the next few steps to create the projects and get them ready for checkout.

First, he creates the MUD directory, RPD_main, where the master RPD will be stored. This master RPD contains the superset of content for the developers. The users will check their projects out of the master, and merge them back in when they want to share their changes. Sally copies her started RPD to the master folder so that Adam can create the first two projects, ProjRevenue and ProjQuota.

First, Adam opens the master RPD in the Administration Tool and selects **Manage > Projects**. Then, in the Project Manager, he selects **Action > New Project**. Adam names the project "ProjRevenue" and proceeds to pick the logical fact tables at the center of the project. The top object in the list expands to show the logical fact tables, but he has a choice of seeing them grouped by the Business Model to which they belong, or by Subject Area.

Figure B-2 shows the different ways Adam can view the logical fact tables.

Figure B-2 Project Dialog with Facts Grouped by Business Model and Subject Area



Adam decides to group facts by Business Model for convenience, although he could have used the Subject Area grouping to select the same fact table. He adds the fact table, plus the default application roles and subject areas specified for this project.

Because there are no custom-defined application roles, users, variables, or initialization blocks yet, he cannot yet add them to the project. Adam repeats this process for ProjQuota, the second project.

Tip: Note that some of the explicit objects are the same in both projects, because both projects share application roles. Similarly, many of the implicit project objects are shared, particularly dimension tables in both the logical and physical models. Keep in mind that projects are a convenience for creating small subsets that are easy to work with; they are not for security. It is critical in your governance process that the owner of each top-level object is assigned and documented for the whole team, because this enables developers to avoid conflicts.

Adam included the logical fact table F10 Bill Rev in the project, even though it is owned by Sally Andre, not by Scott Baker, the owner of this project. He did this because Scott needs to create a measure that derives from measures in both fact tables (Sales percent of quota). Again, the point is to provide the user with the subset of content they need to implement their requirements, not just the objects they own.

Adam saves the master RPD to the shared drive, RPD_Main, as "sales.rpd". It is now ready for users to check out projects and begin working in parallel.

First Developer Checks Out

Now, the two developers will set up their Administration Tool clients for the master repository, check out their projects, and begin working. Sally starts by setting up her Administration Tool client to use the master repository. To do this, she selects **Tools > Options**, and then selects the Multiuser tab. There, she sets up the pointer to the master repository directory. She also enters her full name, which will be useful in logs and locks. Now, she can check out her project and begin working on it.

Meanwhile, in the Master Repository directory, two new files have been created: sales.000 and sales.mhl. [Figure B-3](#) shows the new files.

Figure B-3 Two New Files in the Master Repository Directory

Name	Size	Type	Date Modified
sales.rpd	36 KB	Oracle BI Repository File	6/16/2010 4:39 PM
sales.000	36 KB	000 File	6/16/2010 4:39 PM
sales.mhl	1 KB	MHL File	6/23/2010 3:29 PM

The sales.000 file is an automatic backup created for sales.rpd when Sally checked it out. This file can be used to roll back if problems occur. The sales.mhl file tracks her checkout status and parameters, including project, computer, and user.

Meanwhile, three files have been created in Sally's local repository directory:

- **originalProjRevenue.rpd:** This file is the project subset RPD at the time of checkout. It will be used later as the original in the three-way merge process, and also if Sally discards her changes.
- **ProjRevenue.rpd:** This file contains only the self-consistent subset project (ProjRevenue). This is the file that is open for editing.
- **ProjRevenue.rpd.Log:** This file is the log file for this editing session in the Administration Tool. You can view its contents in the Administration Tool using **File > Multiuser > History**.

[Figure B-4](#) shows the three files in the local repository directory.

Figure B-4 Three New Files in the Local Repository Directory

Name	Size	Type	Date Modified
ProjRevenue.rpd.Log	1 KB	Text Document	6/23/2010 3:29 PM
ProjRevenue.rpd	31 KB	Oracle BI Repositor...	6/23/2010 3:29 PM
originalProjRevenue.rpd	31 KB	Oracle BI Repositor...	6/23/2010 3:29 PM

Now, Sally begins to work on the model for her application. She does not need to change her connection pool settings because she used her own test data source connection pool details when she created the starter content.

Sally starts by opening her fact table and deleting the unused keys based on the modeling best practice. Then, she adds SUM aggregation rules to three measures, Discnt_Value, Revenue, and Units. She also changes the name of Discnt_Value to "Discount Amount," Units to "Units Sold," and Revenue to "Sales Revenue."

Sally also needs to add a new column to the D10 Product table, an upper-case version of the Prod_Dsc column called "PRODUCT DESCRIPTION." It uses the following expression: Upper("Sales"."D10 Product (Dynamic Table)".Prod_Dsc). She also adds dimension hierarchies, creates a variable called "Constant One", and initializes it to the value 1. She uses it to create a new measure, Constant One.

Sally starts her sandbox Oracle Business Intelligence stack so that she can add application roles, and then test her repository using Answers. She follows these steps to start her components in the right order and to configure her system environment:

1. Start the database containing the RCU schema, using its standard controls.
2. Start the Oracle WebLogic Server Administration Server. For example, on Windows, select **Start > Programs > Oracle WebLogic > User Projects > bifoundation_domain > Start Admin Server for WebLogic Server Domain** and enter the user and password created during installation when prompted.

Note that if you used an Enterprise or Software-Only install type, you must also start the Oracle WebLogic Server Managed Server using the Oracle WebLogic Server Administration Console. Typically, you use the Simple install type when installing development sandboxes.

3. Log in to Fusion Middleware Control and upload the repository file, making sure to enter the correct repository password.
4. Also in Fusion Middleware Control, turn off Oracle BI Server caching, so that interpreting the query log is simpler.
5. Still in Fusion Middleware Control, start the system components from the Business Intelligence Overview page.

Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition provides more information about steps 2 - 5.

Because Sally's Oracle BI Server is on a Linux system, she must set up ODBC connectivity on her Windows computer so that her Administration Tool client can access the BI Server there.

Sally manually adds an Oracle BI Server ODBC DSN pointing to the Oracle BI Server on the Linux computer. See "Integrating Other Clients with Oracle Business Intelligence" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for information about how to create an ODBC DSN for the Oracle BI Server.

Sally is using the Oracle WebLogic Server embedded policy store and needs to add two application roles, "Sales Management" and "Sales Rep." To add the roles, she opens a Web browser on her Windows computer and logs in to Fusion Middleware

Control (pointing to her Oracle Business Intelligence stack on Linux). She uses Fusion Middleware Control to create the new roles, maps it to the appropriate users, groups, or other roles, and grants the appropriate permissions to the role.

Tip: See "Creating an Application Role" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Next, Sally needs to add the new application roles to her repository, and then use them for object permissions and data access filters. To accomplish this, Sally does the following:

1. Sally open the Administration Tool and selects **File > Open > Online**. She picks the local Windows ODBC DSN that connects to her Oracle Business Intelligence stack, enters her repository password, and also enters the default user name and password for administering her stack that she created upon install.
2. Next, Sally selects **Manage**, and then selects **Identity** to open the Identity Manager. She clicks **BI Repository** in the navigation tree and then clicks the Application Roles tab. She sees the five default application roles, as well as the new ones she just created.
3. Sally double-clicks the Sales Rep application role, and then clicks **Permissions**. On the Data Filters tab, she adds a data filter with an expression that only allows users who belong to this role to see sales that they themselves have made. On the Object Permissions tab, she sets Read, Read/Write, or No Access permissions that allow Sales Rep users to see revenue, but not quota or cost information. On the Query Limits tab, she keeps the defaults for Max Rows and Max Time, and does not set any time restrictions. She clicks **OK** to return to the Identity Manager.
4. Next, Sally double-clicks the Sales Management application role and sets up Data Filters, Object Permissions, and Query Limits appropriate for this role, based on the decisions of the governance committee.
5. Finally, Sally exits the Identity Manager.

For the new variable and application roles to be in Sally's project the next time she checks it out, she must add them to the project before she checks in her changes. To do this, she performs the same steps that Adam did when he created the projects: She selects **Manage > Projects**, selects her project, selects the new objects in the left pane, and clicks **Add**.

Second Developer Checks Out

While Sally Andre is working on the ProjRevenue project, Scott Baker is getting started on ProjQuota. He set up his Administration Tool options for MUD, checked out his project, and started working.

Scott prefers to work in online mode. Doing this tightens the development/unit test loop, because he is modifying the repository while it is running in the Oracle BI Server. Every time he clicks **Check In Changes** in the Administration Tool toolbar, his changes are applied to the running server. He can then immediately move to Answers and test the changes there. Note that when he adds, deletes, renames, or reorganizes Presentation layer objects, he must reload metadata in the Answers criteria tab to refresh the tree visible there.

First, Scott starts his Oracle Business Intelligence stack, and uploads his checked-out repository using Fusion Middleware Control. He restarts the Oracle BI Server, opens the Administration Tool, and opens his repository in online mode.

Scott must change the connection pool settings to point to his local test database, because the master repository contains Sally's settings. Note that in the merge process,

these connection pool changes will be overridden by the connection pools already in the master repository. Therefore, the next time Scott checks out, he will need to apply his local test connection pool changes again.

Tip: Use the Oracle BI Server XML API to automate connection pool changes required during migrations to production and other environments. See "Moving from Test to Production Environments" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

Scott's next task is to clean up his logical fact table by removing keys. He also gives a measure a SUM aggregation rule and a business-friendly name (Quota Amount).

Scott does not change anything in the F10 logical table because it is owned by Sally. After she checks in, merges, and publishes her changes to the master RPD, he will do the same. Then he will check out again, picking up her changes.

Next, Scott adds a new measure called "Sales percent of quota" to the F50 table. It derives from both fact tables with the following expression:

```
"Sales"."F10 Billed Rev"."Revenue" / "Sales"."F50 Facts Quotas"."Quota Amount"
```

Note that even if Sally changes the name of Revenue in her project, the merge will identify it as the same object and use the new name in Scott's expression. The merge logic can identify the name change because the upgrade ID of the object is still the same as the original.

Finally, Scott forgets what he learned in the Governance Committee meeting, that all the dimensions are owned by Sally. He has a requirement for an all-capitals version of the D10 Product.Prod_Dsc column called PRODUCT DESCRIPTION. He creates a column identical to the one Sally created. This mistake will be detected and resolved through the check-in merge process in a few moments.

Scott does not need to upload his repository and restart his system because he is working in online mode. Instead, he unit tests his changes immediately after making them. Meanwhile, Sally has finished testing her changes.

First Developer Checks In

Sally has finished creating and unit testing her first batch of changes, so she saves her work and prepares to merge it into the master repository. She chooses **File > Multiuser > Merge Local Changes**. If she forgot to add any new objects to a project, a detailed warning is displayed so that she can add the objects to her project and try the merge again. Otherwise, the objects are not extracted the next time she checks out the project.

Next, the Administration Tool locks the master repository so that Sally can merge her changes without any chance of corruption from other users' merges.

Tip: For logging purposes, it is a best practice to use the comment field to provide a description of the changes you are checking in. Checking in frequently also makes it easier to keep track of changes, and easier to audit the history later. Finally, it is a best practice in Administration Tool modeling to work incrementally, which simplifies testing and reduces the complexity of each task.

Sally's changes cause no conflicts, so they do not appear in the Define Merge Strategy step that is displayed next. However, aliases for presentation objects are a special case where you can choose to keep either the modified (your local version) or current (the master), or merge the two. The aliases were automatically created when Sally changed the column names, so that reports written to the old names would not break when she put the new names into production. Because Eden Corporation has no reports yet,

Sally keeps the aliases empty by selecting Current. She does this for "Sales Revenue," "Units Sold," and "Discount Amount."

Tip: Sometimes, there can be a series of aliases if names change more than once. Because there might be a set of reports using the older names, you can select **Merge Choices** in the Define Merge Strategy screen to keep any aliases already in Current as well as the new ones in Modified.

After the merge completes, Sally tests one more time, to ensure that no functionality has been broken by the merge with other content. However, Sally is careful not to test too long, because the lock on the master repository is not released until she publishes the repository to the master directory, or discards her changes.

Finally, Sally selects **File > Multiuser > Publish to Network**. This first copies the master sales.rpd to a backup (sales.001), and then overwrites sales.rpd with the changes from Sally. A merge log is also stored.

Second Developer Checks In

Now that Scott has completed his development work for this phase, he merges his local changes. The Define Merge Strategy screen asks whether to keep the alias created on the presentation column "Quota Amount." Like Sally, Scott chooses to keep the current repository value, which does not use the alias.

After the local merge, Scott unit tests again briefly, although he is aware that he holds the lock on the master repository. Upon inspection, he also notices his mistake of creating the same PRODUCT DESCRIPTION column that Sally did. Because Scott's column was created separately, its internal upgrade ID is different than the one in Sally's. Therefore, even though the name is the same, the merge logic knows it is a different column, and renames it rather than overwriting it by appending #1 (PRODUCT DESCRIPTION#1).

Because Scott has completed the merge, he now has the entire repository open and locked. He deletes the extra column, connects his logic to Sally's PRODUCT DESCRIPTION column, tests again briefly, and publishes his changes to the network master repository.

Note that if Scott had deleted or modified a different user's object, the error might have been more difficult to resolve. It might have required re-creating and equalizing the object, or rolling back to a backup version of the repository and re-creating his own changes.

MUD Administrator Test Migration Activities

To prepare the repository for the test environment, the MUD administrator, Adam Straight, must now perform several tasks directly on the master repository. In other words, he will use **File > Open > Offline** rather than **File > Multiuser > Checkout**.

Adam begins by opening the Administration Tool and then opening sales.rpd in offline mode. As soon as he does this, other users are locked out, and will get Windows permissions errors if they try to check out projects. If Adam needed to open and close the file several times, he would need to remove the RPD from the shared directory while modifying it elsewhere, so that other users would not be able to check out between his changes.

Adam changes the connection pool settings to match the test environment. Note that when Administration Tool users check out projects, connection pool parameters are not included in the checkout. Usually, the master repository in the MUD directory contains the test connection pools, but each individual developer might need different

settings for connecting to their own test databases. At check-in and publish, the connection pools in the master repository are not overwritten by developer changes, so that they can continue to point at the shared test databases.

Adam must also ensure that the new application roles are migrated to the test system. Because there are only two, he decides to reenter them in Fusion Middleware Control on the test system. Adam also provisions some test users or groups to the new application roles so the security filters, permissions, and query limits can be tested.

Finally, Adam uploads the repository to the test system using Fusion Middleware Control and restarts the Oracle BI Server. Using his local Administration Tool, he connects to the test Oracle BI Server in online mode and runs the consistency checker. If any application roles referenced by this repository are missing or incorrect, the consistency checker will list errors for them.

Phase I Testing

The test team can now test the repository. During testing, the test team discovers a bug: "Sales"."F50 Facts Quotas"."Sales percent of quota" was erroneously created with the expression quota/sales instead of sales/quota. The test team writes a bug report, and Scott Baker is assigned to fix the bug.

Scott opens the Administration Tool, checks out ProjQuota, makes the change, changes the connection pool to point to his local test database, and tests on his own sandbox. Then he merges his changes and publishes them to the shared MUD directory. He informs Adam that the bug is fixed and that the repository is ready for him to send to test again.

Adam notes that the connection pools are still pointed at the correct test system, because the MUD feature isolates the master repository from connection pool changes in checked out projects. Adam opens Fusion Middleware Control on the test computer, uploads the repository, and restarts the Oracle BI Server.

The test team tests to completion, and the repository is cleared for production.

Phase I Migration to Production

After the repository has passed the testing phase, it needs its database connection parameters updated and can then be uploaded to production. Also, the application roles must be migrated and provisioned.

Based on the plan provided by the governance team, the production operations team knows the new application roles needed. They create them as Adam did for the test environment. They also provision users or groups to those application roles, based on the security specification from the governance team.

Before migrating to production, Adam has to change the connection pool parameters to the values needed for the production database. In Eden Corporation, Adam has the privilege to see the production connection pools, but the repository developers do not. Therefore, Adam cannot change from the test to production connection pools and leave the repository in the master directory, because the developers have Windows permissions to read and write to it. Instead, he creates an XML patch of the connection pools needed for Production. Then, he copies sales.rpd to a secure directory and applies the patch, and then tests to be sure it really does connect to the production data sources. He then uploads the repository to the production system, and starts the production cluster of servers.

Tip: Use the Oracle BI Server XML API to automate connection pool changes required during migrations to production and other environments. See "Moving from Test to

Production Environments" in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

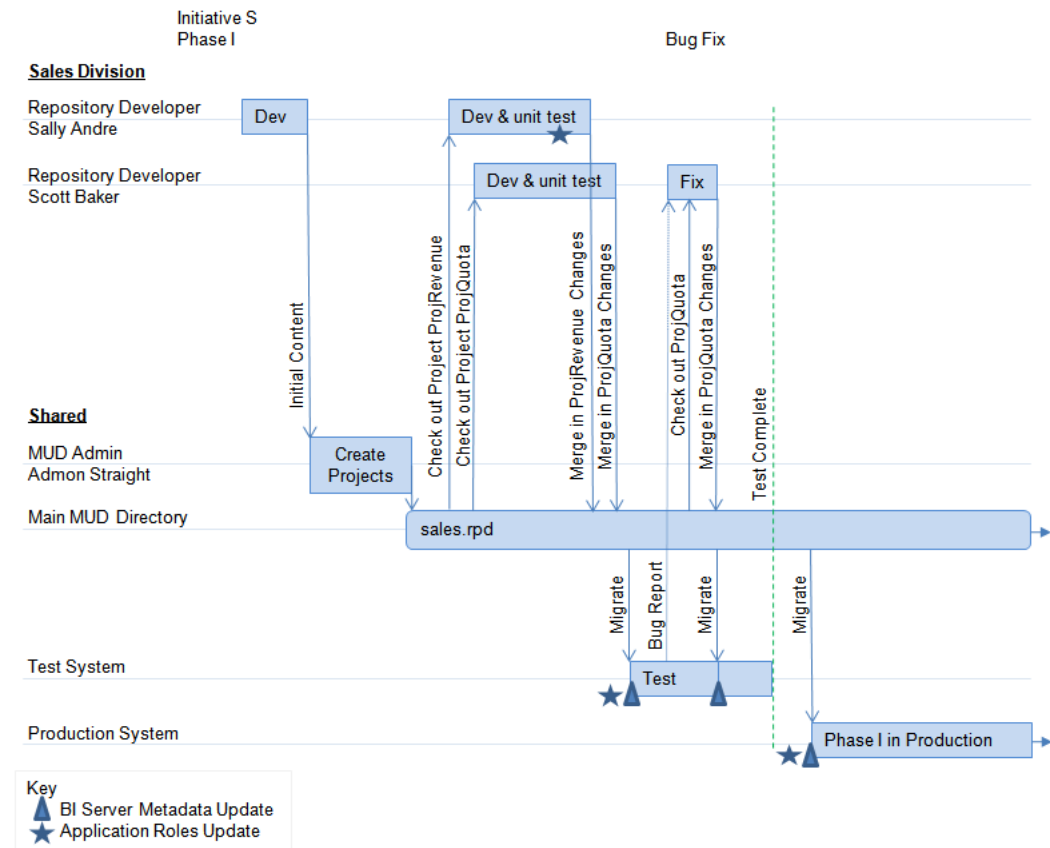
Because the master repository still points to the test databases, the Administration Tool users can still be allowed to see it. Meanwhile, new versions of the production repository can be built at any time by applying the connection pool changes in the XML patch file.

Production validations are now performed. Similar to the migration to the test system, an important validation is to run the consistency checker in online mode to ensure that the application roles are all correct. When this validation is complete, Phase I is in production.

Phase I Summary

Figure B-5 shows the parallel activities for Phase I.

Figure B-5 Summary of Phase I Activities



Phase II - Branching, Fixing, and Patching

In Phase II, development will continue on a new Phase II branch, while a Main branch will track the production application. To manage this work, Adam will add a branch project, and set up a second master repository shared directly, one for Main, and one for the new Phase II branch.

Sally will add more content to ProjRevenue. While she works on that, Scott will add brand new content. After Scott merges and publishes, Adam will create the new

project, ProjTarget, and move Scott's new content into it. Meanwhile, they will have to handle any bugs that occur in production, which is still on the main sales.rpd branch.

The following sections describe Phase II development:

- [Setting Up the Second Branch](#)
- [Developers Check Out Projects](#)
- [Patch Fix for the Main Branch](#)
- [Finishing and Merging Phase II Branch](#)
- [Phase II Summary](#)

Setting Up the Second Branch

Adam begins by creating another MUD directory to hold the master for the new branch. He sets the Windows share security so that Sally and Scott can read or write to it.

Next, Adam places the main repository into the main MUD directory. He adds a new project for the branch, which encompasses all the existing functionality. Then, he closes the repository, and checks out the branch project in his local Administration Tool repository folder. He copies it to the branch MUD directory, where it now serves as the master for the branch.

Developers Check Out Projects

Sally and Scott check out their projects again, and begin developing Sales Initiative Phase II in parallel with each other, and in parallel with Phase I being in production. Because Scott is adding new content that will become a new project, he needs to check out one or more other projects that will provide the shared objects to which he needs to map or join in the new content. He chooses to check out ProjQuota.

Patch Fix for the Main Branch

While Sally and Scott are developing Phase II, an urgent CEO request is escalated to them. The CEO wants the key sales managers to see a new measure called "Sales Quota Variance" on their dashboards within two days.

Scott closes his work on the new project on the Phase II branch; it will stay checked out. Then, he checks out the project that will contain the new measure, ProjQuote, from the main branch master repository (sales.rpd). He creates the new measure and corresponding presentation column, tests it locally, merges it locally, and publishes it back to the main branch.

Scott then reopens the checked-out Phase II repository from his local drive and continues development.

Meanwhile, Adam sends the new sales.rpd to the test environment, where the test team validates the fix.

Next, Adam prepares to send the fixed repository to Production. Rather than send the entire repository, however, he sends a patch of the change.

To create the patch, Adam compares the modified repository to the one that is currently running in production. The repository running in production is the same as the main repository just before the new changes were merged in, so it is one of the backup repositories in the MUD directory. The current repository running in production is the backup called sales.006, the same one he identified as the original for

the upcoming branch merge. He copies this to sales.006.rpd so the Administration Tool can see and open the file. (He cannot simply rename it, because it may be needed for another merge later.)

Figure B-6 shows the files in the MUD directory, including sales.rpd and sales.006.

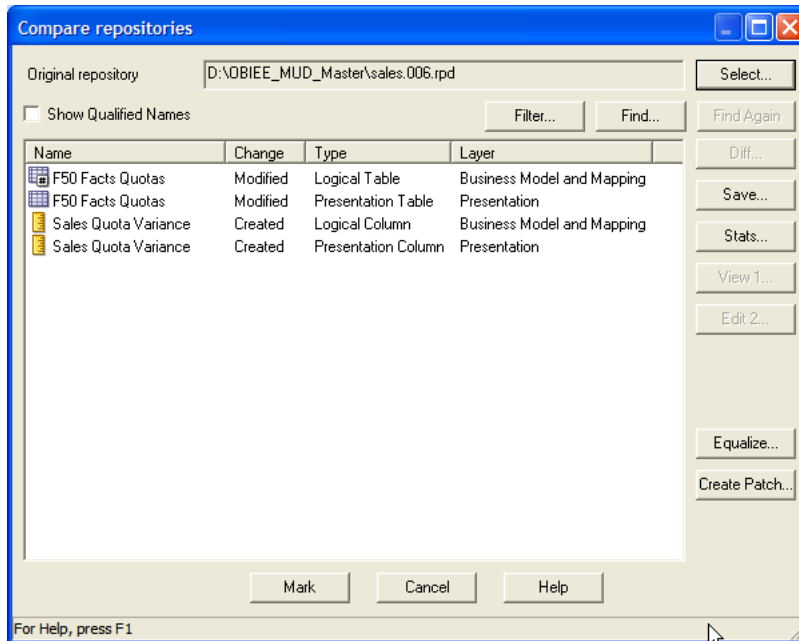
Figure B-6 Renaming sales.006 to sales.006.rpd

Name	Size	Type
sales.000	36 KB	000 File
sales.mhl	5 KB	MHL File
sales.rpd	38 KB	Oracle BI Repositor...
modified subset of sales.001	30 KB	001 File
sales.001	34 KB	001 File
modified subset of sales.002	28 KB	002 File
sales.002	34 KB	002 File
modified subset of sales.003	28 KB	003 File
sales.003	34 KB	003 File
modified subset of sales.004	32 KB	004 File
sales.004	38 KB	004 File
modified subset of sales.005	31 KB	005 File
sales.005	38 KB	005 File
sales.005.csv	1 KB	Microsoft Office Ex...
modified subset of sales.006	31 KB	006 File
sales.006	38 KB	006 File
sales.006.csv	1 KB	Microsoft Office Ex...

Next, Adam opens the repository containing the update, sales.rpd. He selects **File > Compare**, and chooses the sales.006.rpd as the old version to compare. The Compare repositories dialog shows the differences between versions that will be included in the patch.

Figure B-7 shows the Compare repositories dialog.

Figure B-7 Compare Repositories Dialog for sales.rpd and sales.006.rpd



Next, Adam clicks **Create Patch** and saves the result as Patch_variance.xml. The patch contains just the objects needed to apply the two new columns, and their associated interconnections.

Tip: More complex patches might also delete objects, or overwrite objects to merge in new property values.

Adam's patch appears as follows:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<Repository xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <DECLARE>
    <LogicalTable name="F50 Facts Quotas" parentName="&quot;Sales&quot;"
      parentId="2000:68667" parentUid="2160843965" id="2035:69454" uid="2160843966"
      x="718" y="288">
      <Description/>
      <Columns>
        <RefLogicalColumn id="2006:69460" uid="2160844041"
          qualifiedName="&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;.&quot;Quota
          Amount&quot;"/>
        <RefLogicalColumn id="2006:69786" uid="2160845070" qualifiedName=
          "&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          Sales percent of quota&quot;"/>
        <RefLogicalColumn id="2006:70033" uid="2160845342" qualifiedName=
          "&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          Sales Quota Variance&quot;"/>
      </Columns>
      <TableSources>
        <RefLogicalTableSource id="2037:69456" uid="2160844747"
          qualifiedName="&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          F50 Facts Quotas&quot;"/>
      </TableSources>
    </LogicalTable>
    <LogicalColumn name="Sales Quota Variance" parentName=
      "&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;" parentId="2035:69454"
      parentUid="2160843966" id="2006:70033" uid="2160845342" isDerived="true"
      isWriteable="false">
      <Description><![CDATA[quota - sales]]></Description>
      <Expr><![CDATA["Sales"."F50 Facts Quotas"."Quota Amount" - "Sales".
        "F10 Billed Rev"."Sales Revenue" ]]></Expr>
    </LogicalColumn>
    <PresentationTable name="F50 Facts Quotas" parentName=
      "&quot;Sales Quota&quot;.&quot;&quot;"
      parentId="4004:69706" parentUid="2160844968" id="4008:69707"
      uid="2160844969" hasDispName="false" hasDispDescription="false">
      <Description/>
      <Columns>
        <RefPresentationColumn id="4010:69711" uid="2160844973" qualifiedName=
          "&quot;Sales Quota&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          Quota Amount&quot;"/>
        <RefPresentationColumn id="4010:70032" uid="2160845338" qualifiedName=
          "&quot;Sales Quota&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          Sales percent of quota&quot;"/>
        <RefPresentationColumn id="4010:70036" uid="2160845345" qualifiedName=
          "&quot;Sales Quota&quot;.&quot;F50 Facts Quotas&quot;.&quot;
          Sales Quota Variance&quot;"/>
      </Columns>
    </PresentationTable>
    <PresentationColumn name="Sales Quota Variance" parentName="
      &quot;Sales Quota&quot;.&quot;F50 Facts Quotas&quot;" parentId=
      "4008:69707" parentUid="2160844969" id="4010:70036" uid="2160845345"
      hasDispName="false" hasDispDescription="false" overrideLogicalName="false">
      <Description><![CDATA[quota - sales]]></Description>
      <RefLogicalColumn id="2006:70033" uid="2160845342" qualifiedName=
        "&quot;Sales&quot;.&quot;F50 Facts Quotas&quot;.&quot;
        Sales Quota Variance&quot;."/>
    </PresentationColumn>
  </DECLARE>
</Repository>
```

```
        &quot;Sales Quota Variance&quot;"/>  
    </PresentationColumn>  
</DECLARE>  
</Repository>
```

Tip: Unlike migrating an entire repository, there is no need to make any connection pool changes before applying this patch. The correct connection pool settings are already in the repository running in production. The patch will not affect this logic, so the connection pools will stay correct without an intervention.

Finally, Adam must have this patch migrated and applied to the production system. There are several ways to accomplish this:

- 1. Patch main repository offline and upload using Fusion Middleware Control.**
Adam can apply the patch to a copy of the production repository locally on his Windows computer by using the Administration Tool to perform a patch merge. Then, he can upload the repository to the production system using Fusion Middleware Control, like Sally did earlier in her sandbox. Because the production system is clustered, he must restart all the Oracle BI Servers after uploading the repository. Adam can restart manually through Fusion Middleware Control, one server at a time. If he performs a rolling restart in this way, end users do not see any unavailability. Alternatively, Adam or one of the operations staff can write a script using the BI Systems Management API to automate a rolling restart.
- 2. Patch production repository in place using patchrpd utility:** The operations staff can log onto a production system directly, and apply the XML patch using the patchrpd utility. Note that if any conflict occurs, the utility will cancel the update and exit without making changes. If the update is successful, the operations staff can then perform a rolling restart, as described in the previous paragraph.
- 3. Patch running system using biserverxmlcli utility:** This method is not recommended for production systems.

Tip: If you have privileges to log on to a production Oracle BI Server using the Administration Tool in online mode, you can use **File > Copy As** to copy it to your local drive.

Finishing and Merging Phase II Branch

Sally and Scott complete their changes in the new branch, and check them in.

Adam now adds Scott's new content to a new project, projTarget. He performs the same steps as before to send the branch repository to the testing team.

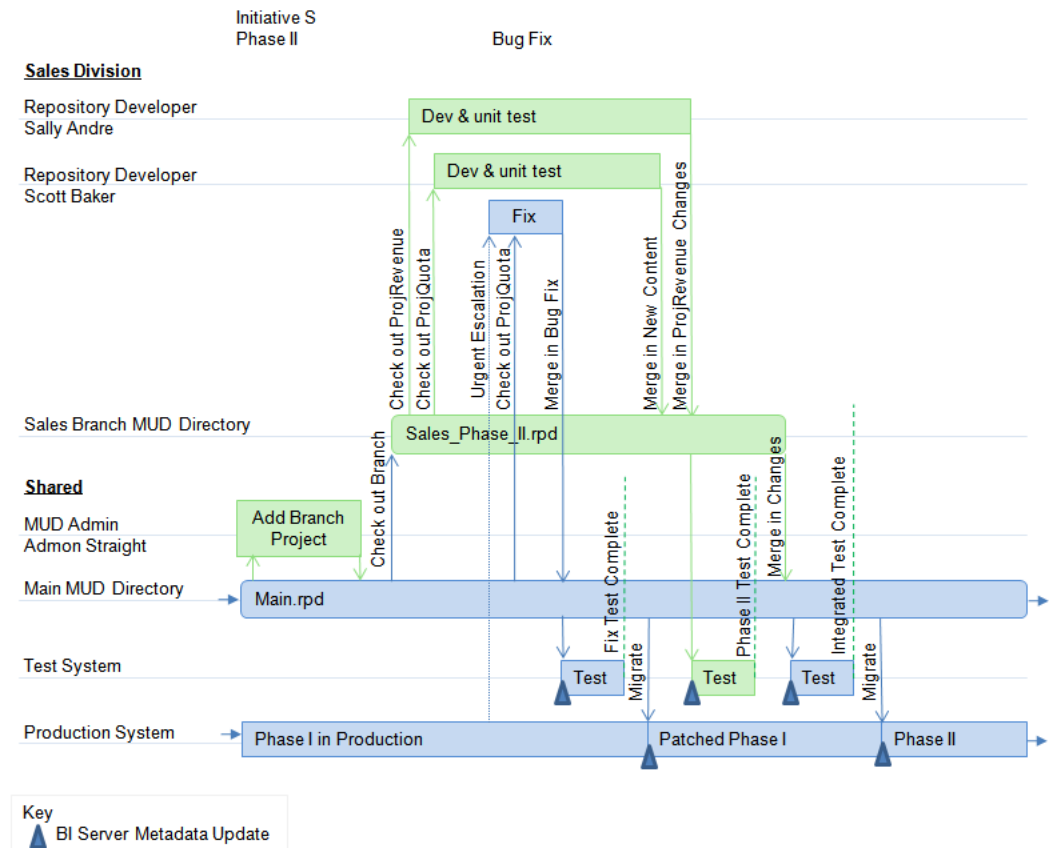
When testing is complete, the branch must be merged back into the main branch using MUD merge. Doing this merges the production patch with the newly developed content, so that can be moved to production later.

Now, sales.rpd contains all the changes, and the branch is no longer needed. Sales.rpd is sent to integrated test, to ensure the merged content does not cause any bugs in the existing content. When integrated testing is complete, Adam creates another patch containing the changes, and has the operations staff apply it to the running production system. Sales Initiative Phase II is now in production.

Phase II Summary

Figure B-8 shows the parallel activities for Phase II.

Figure B-8 Summary of Phase II Activities



Phase III - Independent Semantic Model Development

In the next phase, Sally and Scott begin development of Phase III of the Sales initiative. Meanwhile, Helen Rowe builds the first phase of the HR initiative and brings this new independent semantic model into production.

The following sections describe Phase III development:

- [Security Considerations for Multiple Independent Semantic Models](#)
- [Sales Semantic Model Developers Check Out](#)
- [HR Semantic Model Developer Builds Content](#)
- [Phase III Summary](#)

Security Considerations for Multiple Independent Semantic Models

Helen's application has highly sensitive personal information, such as salaries and medical information. Meanwhile, the Sales application has legally sensitive financial information. Due to corporate security compliance, these two teams are not allowed to see each other's data or metadata. They also have little content they could share, other than generic dimensions like time dimensions. Finally, they have different business drivers, budgets, and schedules.

For these reasons, the Eden Corporation governance committee decided to use independent semantic models in the repository: one for Sales, and the other for HR. This approach requires the two teams to ensure that there are not any shared objects, and there can be no conflicts between their content. The easiest way to ensure this is to

make sure that the names for all top-level objects do not conflict. Even variables and application roles must be different.

Tip: Some governance committees ensure that top-level objects do not conflict by requiring developers to put a prefix specific to each semantic model before the name of each top-level object, such as S_ for Sales and H_ for HR. This practice makes it easy to see which objects belong to which organizations. Other committees prefer to keep a master list of top-level objects, and require new applications to submit top-level object names for review to ensure there are no conflicts. In addition, two-way merges can catch any mistakes before overwrites can damage content or cause unexpected object name changes.

Another security requirement is the need to apply security to the separate MUD directories so that only the correct developers have access to each repository. Sally and Scott can only see and check out from the Sales MUD directory, and Helen can only see and check out from the HR MUD directory. The Main directory continues to exist, since it must hold the merged master that is actually in production, but now only Adam has privileges to see or modify that directory.

At Eden Corporation, a final security requirement is to disable the ability for independent semantic model developers to access the running repository in online mode after the merge. There is only a single repository password, so a developer who has the password and access to the repository can see and modify all its contents in offline mode. However, in online mode, the developer also needs a data access user name and password to log on to the Oracle BI Server. To enforce this security requirement, Adam must ensure that the developers have no privileges to log on to the production or test system in this way. Alternatively, the production operations staff can change the repository password to one that only they know, but this task must be performed on a Windows computer because repository passwords are changed using the Administration Tool.

Sales Semantic Model Developers Check Out

Sally and Scott check out their projects from the new, secure sales branch MUD directory. They begin their work.

HR Semantic Model Developer Builds Content

Because Helen is working alone on her secure, independent semantic model, she does not yet need to check out a project. In fact, she needs to start building her content from a new, blank repository on her local computer. She follows the usual steps of building and unit testing content incrementally.

When she is done with unit testing, Helen has a complete, free-standing repository. She sends it to Adam, who uses a two-way merge to combine it into the main branch repository. Adam performs the following two steps:

1. First, Adam equalizes the two repositories to reassign IDs honoring the different names given to the top-level objects. This practice ensures that there will be no conflicts during the merge.
2. Next, Adam performs a two-way merge by using the Merge Repository Wizard to perform a full merge using a blank repository for the original.

Tip: To create a blank repository, select **File > New Repository**. Then, provide a name (such as blank.rpd) and a repository password. Choose **No** for **Import Metadata** and then click **Finish**.

After the merge, Adam creates a new project for managing the content going forward, hr_payroll. He adds Helen's content to the project. Adam then checks it out of main and posts it to the HR Branch MUD directory. Using a project checkout makes managing IDs and merges easier later.

Adam adjusts connection pool parameters, and migrates the repository to the test computer. When a bug is found, Helen checks out the hr_payroll project, fixes it, unit tests it, and checks it back in. (Note that she checks her functional project out of the checked-out branch project.) Adam migrates it to the test system for further testing. When testing is complete, he checks the completed HR branch repository back into the main branch, and sends the integrated repository to integration testing on the test system.

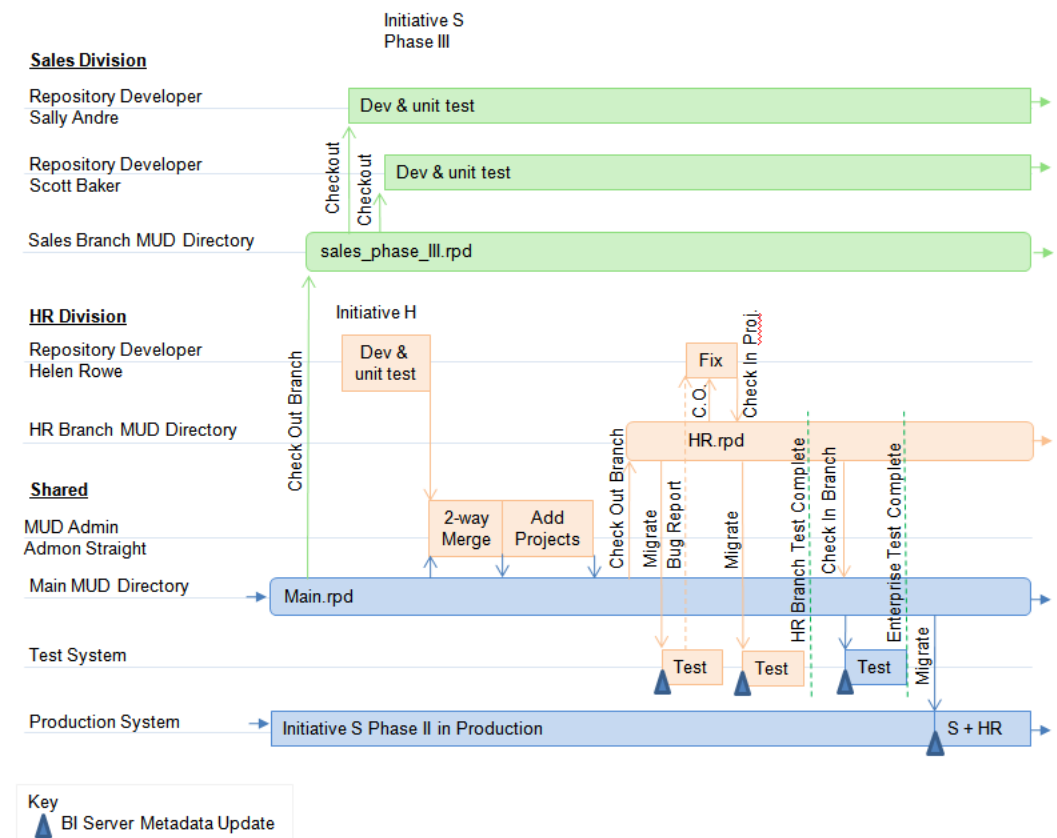
When the integrated repository completes testing, it is ready for migration to production. Again, the options are complete repository migration, or applying a patch to the production environment using patchrpd. Both methods require a rolling restart.

After this step, the production repository contains content for both Initiative S and Initiative H.

Phase III Summary

Figure B-9 shows the parallel activities for Phase III.

Figure B-9 Summary of Phase III Activities



Logical SQL Reference

The Oracle BI Server accepts SQL `SELECT` statements from client tools. Additionally, the Administration Tool enables you to define logical columns with complex expressions. This appendix explains the syntax and semantics for the `SELECT` statement and for the expressions you can use in the Administration Tool to create derived columns.

This reference provides syntax and usage information for the Logical SQL statements understood by the Oracle BI Server. Oracle BI Server Logical SQL includes standard SQL, plus special functions (SQL extensions) like `AGO`, `TODATE`, `EVALUATE`, and others. Logical SQL queries resolve to Presentation layer objects.

The abstraction provided by the Presentation layer and Business Model and Mapping layer enables clients to query data with Logical SQL only, so that the interaction with actual physical sources is handled by the Oracle BI Server. The complexity of the multiple source languages needed to communicate with each data source type is hidden from users and clients.

In Answers, you can view the Logical SQL queries issued by Oracle BI Presentation Services for particular analyses by viewing the SQL Issued section of the Advanced tab of the Analysis editor. If you have the appropriate privileges, then you can also view SQL by displaying the Manage Sessions page in the Administration tab. Click **View Log** from the Manage Sessions page to see further details.

In Answers, there are also several places where you can issue Logical SQL. If you have the appropriate privileges, then you can use the Issue SQL page in the Administration tab to enter any SQL code to send to the Oracle BI Server. If an analysis does not contain hierarchical columns, member selections, or groups, then you can use the Advanced SQL Clauses fields in the Advanced tab of the Analysis editor. You can also enter SQL in the New Filter dialog.

In the Administration Tool, Logical SQL appears mostly in the form of expressions related to objects in the Business Model and Mapping layer. You typically create SQL functions in Expression Builder; see "[About the Expression Builder Dialogs](#)" for a summary of the places in the Administration Tool where you can build Logical SQL expressions.

Other clients, like Oracle BI Publisher, Oracle's Hyperion Interactive Reporting, Smart View, the Oracle BI Add-in for Microsoft Office, and Essbase, also provide their own interfaces to view and issue Logical SQL to the Oracle BI Server.

This appendix contains the following topics:

- [SQL Syntax and Semantics](#)
- [Aggregate, Running Aggregate, and Time Series Functions](#)
- [String Functions](#)

- [Math Functions](#)
- [Calendar Date/Time Functions](#)
- [Conversion Functions](#)
- [Database Functions](#)
- [Hierarchy Navigation Functions](#)
- [System Functions](#)

SQL Syntax and Semantics

This section explains SQL syntax and semantics. The following topics are included:

- [Syntax and Usage Notes for the SELECT Statement](#)
- [Syntax and Usage Notes for SELECT_PHYSICAL](#)
- [Rules for Queries with Aggregate Functions](#)
- [Operators](#)
- [Conditional Expressions](#)
- [Expressing Literals](#)
- [Calculated Members](#)
- [Variables](#)

Syntax and Usage Notes for the SELECT Statement

The `SELECT` statement, or query specification, is the way to query a decision support system through the Oracle BI Server. A `SELECT` statement returns a table to the client that matches the query. It is a table in the sense that the results are in the form of rows and columns.

The `SELECT` statement is the basis for querying any structured query language (SQL) database. The Oracle BI Server accepts logical requests to query objects in a repository, and users (or query tools) make those logical requests with ordinary SQL `SELECT` statements. The server then translates the logical requests into physical queries against one or more data sources, combines the results to match the logical request, and returns the answer to the end user.

The `SELECT` statement in Logical SQL differs from standard SQL in that tables do not need to be joined. Any join conditions supplied in the query are ignored because the join conditions are predefined in the Oracle BI repository.

This section provides the basic syntax for the `SELECT` statement, as well as definitions for individual clauses. The syntax descriptions cover only basic syntax and features unique to the Oracle BI Server. For a more comprehensive description of SQL syntax, see a third-party reference book on SQL or a reference manual on SQL from your database vendors. For Oracle Database, see *Oracle Database SQL Language Reference*.

This section contains the following topics:

- [Basic Syntax for the SELECT Statement](#)
- [Usage Notes](#)
- [Subquery Support](#)
- [SELECT List Syntax](#)

- [FROM Clause Syntax](#)
- [WHERE Clause Syntax](#)
- [GROUP BY Clause Syntax](#)
- [ORDER BY Clause Syntax](#)

Basic Syntax for the SELECT Statement

Syntax for the `SELECT` statement is as follows:

```
SELECT [DISTINCT] select_list
FROM from_clause
[WHERE search_condition]
[GROUP BY column {, column}
  [HAVING search_condition]]
[ORDER BY column {, column}]
```

Where:

`select_list` is the list of columns specified in the request. See "[SELECT List Syntax](#)" for more information.

`FROM from_clause` is the list of tables in the request. Optionally includes certain join information for the request. See "[FROM Clause Syntax](#)" for more information.

`WHERE search_condition` specifies any combination of conditions to form a conditional test. A `WHERE` clause acts as a filter that lets you constrain a request to obtain results that answer a particular question. Together with the columns you select, filters determine what your results will contain. See "[WHERE Clause Syntax](#)" for more information.

`GROUP BY column {, column}` specifies a column (or alias) belonging to a table defined in the data source. See for more information.

`HAVING search_condition` specifies any combination of conditions to form a conditional test. The syntax is identical to that for the `WHERE` clause.

`ORDER BY column {, column}` specifies the columns to order the results by. See "[ORDER BY Clause Syntax](#)" for more information.

Usage Notes

The Oracle BI Server treats the `SELECT` statement as a logical request. If aggregated data is requested in the `SELECT` statement, a `GROUP BY` clause is automatically assumed by the server. Any join conditions supplied in the query are ignored because the join conditions are all predefined in the Oracle BI repository.

The Oracle BI Server accepts the following SQL syntaxes for comments:

- `/**/` C-style comments
- `//` Double slash for single-line comments
- `#` Number sign for single-line comments

Subquery Support

The Oracle BI Server supports certain subqueries, as well as `UNION`, `UNION ALL`, `INTERSECT`, and `EXCEPT` operations in logical requests. This functionality increases the range of business questions that can be answered, eases the formulation of queries, and provides some ability to query across multiple business models.

The Oracle BI Server supports the following subquery predicates in any conditional expression (for example, within `WHERE`, `HAVING`, or `CASE` statements):

```
IN, NOT IN
Any, >=Any, =Any, <Any, <=Any, <>Any
All, >=All, =All, <All, <=All, <>All
EXISTS, NOT EXISTS
```

In Answers, advanced users and developers can use the Advanced SQL Clauses fields in the Advanced tab of the Analysis editor to specify various SQL clauses, such as `GROUP BY`, `HAVING`, and `DISTINCT`, to include in the SQL queries that are sent to the Oracle BI Server. If an analysis contains hierarchical columns, selections, or groups, then certain Advanced SQL Clauses fields are not available.

SELECT List Syntax

The `select_list` lists the columns in the request. All columns need to be from a single business model. Table names can be included (as `Table.Column`), but are optional unless column names are not unique within a business model. If column names contain spaces, enclose column names in double quotes. The `DISTINCT` keyword does not need to be included, because the Oracle BI Server always does a distinct query. Columns that are being aggregated do not need to include the aggregation function (such as `SUM`), as aggregation rules are known to the server and aggregation is performed automatically.

Syntax

```
...
* |
  (column | expr) [[AS] alias]
  {, (column | expr) [[AS] alias] }
...
```

Where:

* Indicates all columns in the resultant table in the `FROM` clause.

`column` is a column (or alias) belonging to a table defined in the data source.

`expr` is any valid SQL expression.

FROM Clause Syntax

The Oracle BI Server accepts any valid SQL `FROM` clause syntax. To simplify `FROM` clause creation, you can specify the name of a subject area instead of a list of tables. The Oracle BI Server determines the proper tables and the proper join specifications based on the columns the request asks for and the configuration of the Oracle BI repository.

WHERE Clause Syntax

The Oracle BI Server accepts any valid SQL `WHERE` clause syntax. There is no need to specify any join conditions in the `WHERE` clause, because the joins are all configured within the Oracle BI repository. Any join conditions specified in the `WHERE` clause are ignored.

The Oracle BI Server also supports the following subquery predicates in any conditional expression (`WHERE`, `HAVING` or `CASE` statements):

```
IN, NOT IN
Any, >=Any, =Any, <Any, <=Any, <>Any
```


All, >=All, =All, <All, <=All, <>All
 EXISTS, NOT EXISTS

GROUP BY Clause Syntax

With auto aggregation on the Oracle BI Server, there is no need to submit a GROUP BY clause. When no GROUP BY clause is specified, the GROUP BY specification defaults to all of the nonaggregation columns in the SELECT list. If you explicitly use aggregation functions in the select list, you can specify a GROUP BY clause with different columns and the Oracle BI Server computes the results based on the level specified in the GROUP BY clause.

See ["Rules for Queries with Aggregate Functions"](#) for additional details, as well as some examples of using the GROUP BY clause in queries against the Oracle BI Server.

ORDER BY Clause Syntax

The Oracle BI Server accepts any valid SQL ORDER BY clause syntax, including referencing columns by their order in the select list (such as ORDER BY 3, 1, 5).

In addition, you can use the following syntax to alter the sort order for nulls in the query:

```
ORDER BY col1 NULLS LAST, ORDER BY col2 NULLS FIRST
```

Syntax and Usage Notes for SELECT_PHYSICAL

The SELECT_PHYSICAL command provides the functionality to directly query objects in the Physical layer of the metadata repository, and to nest such a statement within a query against the Business Model and Mapping layer or the Presentation layer.

Though a SELECT_PHYSICAL query bypasses the Presentation layer and the Business Model and Mapping layer, the Oracle BI Server still performs parsing, interpretation, and query generation on a SELECT_PHYSICAL query before passing it to the database.

A SELECT_PHYSICAL command can contain any element allowed in standard Oracle BI Server SQL with the following constraints:

- The SELECT_PHYSICAL command does not explicitly reference structures in the repository Business Model and Mapping layer or the Presentation layer
- The SELECT_PHYSICAL command does not require implicit logical transformation
- The SELECT_PHYSICAL command cannot contain certain aggregate functions - see ["Aggregate Functions Not Supported in SELECT_PHYSICAL Queries"](#) for details

Note: SELECT_PHYSICAL statements are not cached.

You can set up an ODBC connection to the Oracle BI Server to be a dedicated physical connection over which all SELECT queries are treated as SELECT_PHYSICAL queries. To do this, select **Route Requests To Physical Layer** in the ODBC data source for the Oracle BI Server. See *"Integrating Other Clients with Oracle Business Intelligence"* in *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition* for more information.

SELECT_PHYSICAL statements are logged as Physical Request entries.

The topics in this section are the following:

- [Syntax for the SELECT_PHYSICAL Statement](#)
- [Aggregate Functions Not Supported in SELECT_PHYSICAL Queries](#)
- [Queries Supported by SELECT_PHYSICAL](#)
- [Using the NATURAL_JOIN Keyword](#)
- [Special Usages of SELECT_PHYSICAL](#)

Syntax for the SELECT_PHYSICAL Statement

Basic syntax for SELECT_PHYSICAL queries is equivalent to "[Basic Syntax for the SELECT Statement](#)" with the term SELECT_PHYSICAL replacing the word SELECT, namely:

```
SELECT_PHYSICAL [DISTINCT] select_list
FROM from_clause
[WHERE search_condition]
[GROUP BY column {, column}
   [HAVING search_condition]]
[ORDER BY column {, column}]
```

Notes: The SELECT_PHYSICAL statement is close to the standard ANSI SQL SELECT statement. For example, you cannot omit the GROUP BY clause nor, where relevant, the HAVING clause in a SELECT_PHYSICAL aggregate query.

In SELECT_PHYSICAL queries, you must fully qualify the table names in the FROM list. Each fully qualified table name must match a table name in the physical layer of the repository.

A fully qualified table name consists of up to four components, database name, catalog name, schema name, and table name. Each component is surrounded by double quotes (") with a period (.) separator between components. For example, "SQL_DB"."My_Catalog"."My_Schema"."Customers" for a SQL Server table, and "FoodMart"... "Sales" for a cube table.

Refer to the corresponding topics in "[Basic Syntax for the SELECT Statement](#)" for more information about the different clauses and sub-clauses of the SELECT_PHYSICAL command.

Aggregate Functions Not Supported in SELECT_PHYSICAL Queries

The following aggregate functions are not supported in SELECT_PHYSICAL queries:

- AGO
- BOTTOMN
- FILTER
- FIRST
- LAST
- RCOUNT
- RMAX
- RMIN

- RSUM
- TODATE
- TOPN

Queries Supported by SELECT_PHYSICAL

The Oracle BI Server supports the use of `SELECT_PHYSICAL` for the following types of logical query:

- **Standard Non-Aggregate Queries**

Standard non-aggregate `SELECT_PHYSICAL` commands follow the same rules as standard non-aggregate `SELECT` commands. They can also include scalar functions, such as String, Math, and Calendar Date/Time functions. For example:

```
SELECT_PHYSICAL productid, categoryid
FROM "My_DB"."My_Schema"."products"
WHERE categoryid > 5;
```

```
SELECT_PHYSICAL LEFT(productname,10)
FROM "My_DB"."My_Schema"."products"
WHERE productname is not null;
```

- **Queries with Aggregate Functions**

In general, all aggregate functions supported in `SELECT` queries are also supported in `SELECT_PHYSICAL` queries. See ["Aggregate Functions Not Supported in SELECT_PHYSICAL Queries"](#) for a list of the exceptions to this rule.

For aggregates supported in `SELECT_PHYSICAL` commands, each aggregate must have an explicitly specified aggregation level, using the `GROUP BY` clause or the `BY` clause. For example:

```
SELECT_PHYSICAL employeeid, SUM(quantity by)
FROM "My_DB"."My_Schema"."employees";
```

```
SELECT_PHYSICAL employeeid, SUM(quantity)
FROM "My_DB"."My_Schema"."employees"
GROUP BY employeeid
HAVING SUM(quantity) > 100;
```

- **Subqueries**

The Oracle BI Server supports the following types of query:

- Queries where both the parent query and the subquery use `SELECT_PHYSICAL`
- Parent query uses `SELECT` and subquery uses `SELECT_PHYSICAL`

Subqueries are supported on both filters and on projections embedded in a Case statement.

For example:

```
SELECT_PHYSICAL *
FROM "My_DB"."My_Schema"."products"
WHERE supplierid IN
  (SELECT_PHYSICAL supplierid
   FROM "My_DB"."My_Schema"."suppliers");
```

```
SELECT productid
FROM snowflakesales.product
```

```
WHERE categoryid IN
  (SELECT_PHYSICAL categoryid
   FROM "My_DB"."My_Schema"."categories");

SELECT CASE WHEN b.categoryid IN
  (SELECT_PHYSICAL a.categoryid
   FROM "My_DB"."My_Schema"."products" a)
  THEN b.categoryid END
FROM categories b;
```

■ Queries with Derived Tables

Both `SELECT` and `SELECT_PHYSICAL` queries can have derived tables in their `FROM` clause. The tables can be derived using either `SELECT` or `SELECT_PHYSICAL`. For example:

```
SELECT_PHYSICAL COUNT(DISTINCT t.rto)
FROM
  (SELECT_PHYSICAL employeed AS id, reportsto AS rto
   FROM "My_DB"."My_Schema"."employees") t;

SELECT productid, categoryid
FROM
  (SELECT_PHYSICAL productid, categoryid
   FROM "My_DB"."My_Schema"."products" a
   LEFT OUTER JOIN "My_DB"."My_Schema"."categories" b
   ON a.categoryid = b.categoryid);

SELECT y.cid, sum(x.qty)
FROM
  (SELECT productid pid, categoryid cid, qtysold qty
   FROM sales.product) x
  RIGHT OUTER JOIN
  (SELECT_PHYSICAL CASE categoryid WHEN 1 THEN null ELSE categoryid END cid
   FROM "My_DB"."My_Schema"."categories") y
  ON x.cid = y.cid
  GROUP BY y.cid;
```

■ Cross-Database Queries

You can use `SELECT_PHYSICAL` to join tables in different databases. For example:

```
SELECT_PHYSICAL a.productid, b.categoryid
FROM "My_DB"."My_Schema"."products" a
  FULL OUTER JOIN
  "My_DB2"."My_Schema"."categories" b
  ON a.categoryid = b.categoryid
```

Using the `NATURAL_JOIN` Keyword

`SELECT_PHYSICAL` queries support the `NATURAL JOIN` syntax, which enables you to use predefined join expressions. For ADF Business Component data sources, the ViewLink in ADF becomes active. The `NATURAL JOIN` join type, however, is not exposed for use in Logical Table Sources (for example, `LEFT OUTER JOIN`).

You can only use the `NATURAL JOIN` keyword in `SELECT_PHYSICAL` queries. The `NATURAL JOIN` behavior in Oracle Business Intelligence is different from the ANSI `NATURAL JOIN`. The following examples illustrate how joins are executed with and without the `NATURAL JOIN` syntax:

```
SELECT PHYSICAL *
FROM A, B;
```

In this example, no join is executed between A and B (even if one is defined in the metadata).

```
SELECT_PHYSICAL *
FROM A NATURAL JOIN B;
```

In this example, the physical join between A and B is executed. For ADF Business Component data sources, the join expression defined by the underlying ViewLink is used.

```
SELECT_PHYSICAL *
FROM C, A NATURAL JOIN B;
```

In this example, even if C is joined to A in the metadata, only the A-B join is active. The C-A join is not used.

Special Usages of SELECT_PHYSICAL

You can use session variables and the INDEXCOL function in a SELECT_PHYSICAL command, as in the following examples:

```
SELECT_PHYSICAL VALUEOF(NQ_SESSION.REGION)
FROM "My_DB"."My_Schema"."products";
```

```
SELECT_PHYSICAL INDEXCOL(VALUEOF(NQ_SESSION.INDEXCOLINDEX), productid, categoryid)
FROM "My_DB"."My_Schema"."products";
```

Rules for Queries with Aggregate Functions

The Oracle BI Server simplifies the SQL statements needed to craft aggregate queries. This section outlines the rules that the Oracle BI Server follows for whether a query contains a GROUP BY clause and, if a GROUP BY clause is specified, what results you should expect from the query. The rules outlined in this section apply to all aggregates used in SQL statements (SUM, AVG, MIN, MAX, COUNT (*), and COUNT).

Computing Aggregates of Baseline Columns

A baseline column is a column that has no aggregation rule defined in the Aggregation tab of the Logical Column dialog in the repository. Baseline columns map to nonaggregated data at the level of granularity of the logical table to which they belong. If you perform aggregation (SUM, AVG, MIN, MAX, or COUNT) on a baseline column through a SQL request, the Oracle BI Server calculates the aggregation at the level based on the following rules:

- If there is no GROUP BY clause specified, the level of aggregation is grouped by all of the nonaggregate columns in the SELECT list.
- If there is a GROUP BY clause specified, the level of aggregation is based on the columns specified in the GROUP BY clause.

For example, consider the following query, where the column revenue is defined in the repository as a baseline column (no aggregation rules specified in the Logical Column > Aggregation tab):

```
SELECT year, product, SUM(revenue)
FROM time, products, facts
```

The results appear in the following list by year, products, and then sum of revenue.

YEAR	PRODUCT	SUM(REVENUE)
1998	Coke	500
1998	Pepsi	600
1999	Coke	600
1999	Pepsi	550
2000	Coke	800
2000	Pepsi	600

This query returns results grouped by year and product, or in other words, it returns one row for each product and year combination. The sum calculated for each row is the sum of all the sales for that product in that year. It is logically the same query as the following:

```
SELECT year, product, SUM(revenue)
FROM time, products, facts
GROUP BY year, product
```

If you change the `GROUP BY` clause to only group by year, then the sum calculated is the sum of all products for the year, as follows:

```
SELECT year, product, SUM(revenue)
FROM time, products, facts
GROUP BY year
```

The results appear in the following list by year, products, and then sum of revenue:

YEAR	PRODUCT	SUM(REVENUE)
1998	Coke	1100
1998	Pepsi	1100
1999	Coke	1150
1999	Pepsi	1150
2000	Coke	1400
2000	Pepsi	1400

If you add a column to the query requesting the `COUNT` of revenue, the Oracle BI Server calculates the number of records used to calculate the results for each group. In this case, it is a year, as shown in the following example:

```
SELECT year, product, SUM(revenue), COUNT(revenue)
FROM time, products, facts
GROUP BY year
```

The results appear in the following list by year, products, sum of revenue, and then revenue count:

YEAR	PRODUCT	SUM(REVENUE)	COUNT(REVENUE)
1998	Coke	1100	6000
1998	Pepsi	1100	6000
1999	Coke	1150	6500

YEAR	PRODUCT	SUM(REVENUE)	COUNT(REVENUE)
1999	Pepsi	1150	6500
2000	Coke	1400	8000
2000	Pepsi	1400	8000

Computing Aggregates of Measure Columns

A measure column is a column that has a default aggregation rule defined in the Aggregation tab of the Logical Column dialog in the repository. Measure columns always calculate the aggregation with which they are defined. If you perform explicit aggregation (SUM, AVG, MIN, MAX, or COUNT) on a measure column through a SQL request, you are actually asking for an aggregate of an aggregate. For these nested aggregates, the Oracle BI Server calculates the aggregation based on the following rules:

- A request for a measure column without an aggregate function defined in a SQL statement is always grouped at the level of the nonaggregate columns in the SELECT list, regardless of whether the query specifies a GROUP BY clause.
- If there is no GROUP BY clause specified, the nested aggregate is a grand total of each group determined by all of the nonaggregate columns in the SELECT list.
- If there is a GROUP BY clause specified, the nested aggregation calculates the total for each group as specified in the GROUP BY clause.

For example, consider the following query, where the column SumOfRevenue is defined in the repository as a measure column with a default aggregation rule of SUM (SUM aggregation rule specified in the Aggregation tab of the Logical Column dialog):

```
SELECT year, product, SumOfRevenue, SUM(SumOfRevenue)
FROM time, products, facts
```

The following query results are grouped by year and product, or in other words, it returns one row for each product and year combination. The sum calculated for each row in the SumOfRevenue column is the sum of all the sales for that product in that year because the measure column is always at the level defined by the nonaggregation columns in the query.

YEAR	PRODUCT	SUMofREVENUE	SUM(SUMofREVENUE)
1998	Coke	500	3650
1998	Pepsi	600	3650
1999	Coke	600	3650
1999	Pepsi	550	3650
2000	Coke	800	3650
2000	Pepsi	600	3650

If you set the GROUP BY clause to only group by year, then the sum calculated in the SumOfRevenue column is the sum of each product for the year, and the sum calculated in the SUM(SumOfRevenue) column is total sales of all products for the given year. The following is the query:

```
SELECT year, product, SumOfRevenue, SUM(SumOfRevenue)
FROM time, products, facts
GROUP BY year
```

In the following result set, the sum calculated for each row in the `SumOfRevenue` column is the sum of all the sales for that product in that year because the measure column is always at the level defined by the nonaggregation columns in the query. The `SUM(SumOfRevenue)` is the same for each row corresponding to a given year, and that sum represents the total sales for that year. In this case, it is the sales of Coke plus the sales of Pepsi.

YEAR	PRODUCT	SUMofREVENUE	SUM(SUMofREVENUE)
1998	Coke	500	1100
1998	Pepsi	600	1100
1999	Coke	600	1150
1999	Pepsi	550	1150
2000	Coke	800	1400
2000	Pepsi	600	1400

Display Function Reset Behavior

A display function is a function that operates on the result set of a query. The display functions the Oracle BI Server supports (`RANK`, `TOPN`, `BOTTOMN`, `PERCENTILE`, `NTILE`, `MAVG`, `MEDIAN`, and varieties of standard deviation) are specified in the `SELECT` list of a SQL query. Queries that use display functions conform to the following rules:

- If no `GROUP BY` clause is specified, the display function operates across the entire result set, or in other words, the grouping level for the display function matches that of the query.
- If there is a `GROUP BY` clause specified, the display function resets its values for each group as specified in the `GROUP BY` clause.

For example, in the following query, `SumOfRevenue` is defined as a measure column with the default aggregation rule of `SUM`:

```
SELECT year, product, SumOfRevenue, RANK(SumOfRevenue)
FROM time, products, facts
```

In the following query result set, there is no `GROUP BY` clause specified, so the rank is calculated across the entire result set:

YEAR	PRODUCT	SUMofREVENUE	RANK(SUMofREVENUE)
1998	Coke	500	6
1998	Pepsi	600	2
1999	Coke	600	2
1999	Pepsi	550	5
2000	Coke	800	1
2000	Pepsi	600	2

If you change the `GROUP BY` clause to group by year and product, then the rank is reset for each year, as follows:

```
SELECT year, product, SUM(revenue), RANK(sum(revenue) by year)
FROM time, products, facts
```



```
GROUP BY year, product
```

In the following result set, the rank is reset each time the year changes, and because there are two rows for each year, the value of the rank is always 1 or 2:

YEAR	PRODUCT	SUMofREVENUE	RANK(SUM(REVENUE) by year)
1998	Coke	500	2
1998	Pepsi	600	1
1999	Coke	600	1
1999	Pepsi	550	2
2000	Coke	800	1
2000	Pepsi	600	2

Alternative Syntax

When using an aggregate function, you can calculate a specified level of aggregation using `BY` within the aggregate function. If you do this, you do not need a `GROUP BY` clause.

For example, the following query returns the column `year_revenue` that displays revenue aggregated by year:

```
SELECT year, product, revenue, SUM(revenue BY year) as year_revenue
FROM softdrinks
```

The same syntax can be used with display functions. The following query calculates overall rank of revenue for each product for each year (each row in the entire result set), and also the rank of each product's revenue within each year:

```
SELECT year, product, revenue, rank(revenue), RANK(revenue by year)
FROM softdrinks ORDER BY 1, 5
```

Using FILTER to Compute a Conditional Aggregate

In SQL query language, traditional aggregates, such as `SUM`, `COUNT`, `MIN`, and `MAX` are evaluated on a group of tuples (an ordered list of objects, each of a specified type), determined by the `GROUP BY` clause. All the aggregates specified in the `SELECT` clause of a query are evaluated over the same subset of tuples. Conditional aggregates extend SQL by restricting their input using a predicate.

`FILTER` is an operator that restricts the set of rows used to compute its aggregate argument to rows that satisfy the `USING` condition. The `FILTER` operator is a Logical SQL construct. It may be used in logical queries referring to the metadata, or in logical columns that use existing logical columns as the source.

Syntax

Conditional aggregates are only notational concepts and they do not represent executable operators. Conditional aggregates are expressed in the form of a function as shown in the following statement:

```
FILTER(measure_expr USING boolean_expr)
```

Where:

measure_expr is an expression that contains at least one measure. The following is a list of examples:

- The expression `Sales + 1` is allowed if `Sales` is a measure.
- The expression `productid` is not allowed if `productid` is a scalar attribute.

boolean_expr is a boolean expression (evaluates to `TRUE` or `FALSE`) that does not contain any measures. This expression may not contain any nested queries.

Example

The following is a simple example of the `FILTER` function:

```
SELECT year,
FILTER(sales USING product = 'coke'),
FILTER(sales USING product = 'pepsi')
FROM logBeverages
```

After navigation, this query is executed as follows:

```
SELECT year,
SUM(CASE WHEN product = 'coke' THEN sales),
SUM(CASE WHEN product = 'pepsi' THEN sales)
FROM physBeverages
WHERE product = 'coke' OR product = 'pepsi'
GROUP BY year
```

Error Handling

In the example `FILTER(x USING y)`, error messages are returned in the following situations:

- The *y* expression is not a boolean expression.
- The *y* expression contains measures.
- `FILTER` is used in outer query block.
- Explicit aggregates are used in the *x* (measure) expression. For example, `FILTER(COUNT(product), C)`.

Operators

There are two types of operators: SQL logical operators, and mathematical operators.

SQL Logical Operators

The following SQL logical operators are used to specify comparisons between expressions.

- **Between:** Used to determine boundaries for a condition. Each boundary is an expression, and the bounds do not include the boundary limits, as in less than and greater than (as opposed to less than or equal to and greater than or equal to). `BETWEEN` can be preceded with `NOT` to negate the condition.
- **In:** Specifies a comparison of a column value with a set of values.
- **Is Null:** Specifies a comparison of a column value with the null value.
- **Like:** Specifies a comparison to a literal value. Often used with wildcard characters to indicate any character string match of zero or more characters (%) or a any single character match (_).

Mathematical Operators

Mathematical operators are used to combine expression elements to make certain types of comparisons in an expression.

Table C-1 lists operators and describes their use in an expression.

Table C-1 Operators

Operator	Description
+	Plus sign for addition.
-	Minus sign for subtraction.
*	Multiply sign for multiplication.
/	Divide by sign for division.
	Character string concatenation.
(Open parenthesis.
)	Closed parenthesis.
>	Greater than sign, indicating values higher than the comparison.
<	Less than sign, indicating values lower than the comparison.
=	Equal sign, indicating the same value.
<=	Less than or equal to sign, indicating values the same or lower than the comparison.
>=	Greater than or equal to sign, indicating values the same or higher than the comparison.
<>	Not equal to, indicating values higher or lower, but different.
AND	AND connective, indicating intersection with one or more conditions to form a compound condition.
OR	OR connective, indicating the union with one or more conditions to form a compound condition.
NOT	NOT connective, indicating a condition is not met.
,	Comma, used to separate elements in a list.

Conditional Expressions

Expressions are building blocks for creating conditional expressions that convert a value from one form to another. Expressions include:

- [CASE \(Switch\)](#)
- [CASE \(If\)](#)

CASE (Switch)

This form of the CASE statement is also referred to as the CASE (Lookup) form. The value of *expr1* is examined, then the WHEN expressions. If *expr1* matches any WHEN expression, it assigns the value in the corresponding THEN expression.

If none of the WHEN expressions match, it assigns the default value specified in the ELSE expression. If no ELSE expression is specified, the system automatically adds an ELSE NULL.

If *expr1* matches an expression in multiple WHEN clauses, only the expression following the first match is assigned.

Note: In a CASE statement, AND has precedence over OR.

Syntax

```
CASE expr1
  WHEN expr2 THEN expr3
  {WHEN expr... THEN expr...}
  ELSE expr
END
```

Where:

CASE starts the CASE statement. Must be followed by an expression and one or more WHEN and THEN statements, an optional ELSE statement, and the END keyword.

WHEN specifies the condition to be satisfied.

THEN specifies the value to assign if the corresponding WHEN expression is satisfied.

ELSE specifies the value to assign if none of the WHEN conditions are satisfied. If omitted, ELSE NULL is assumed.

END ends the CASE statement.

Example

```
CASE Score-par
  WHEN -5 THEN 'Birdie on Par 6'
  WHEN -4 THEN 'Must be Tiger'
  WHEN -3 THEN 'Three under par'
  WHEN -2 THEN 'Two under par'
  WHEN -1 THEN 'Birdie'
  WHEN 0 THEN 'Par'
  WHEN 1 THEN 'Bogey'
  WHEN 2 THEN 'Double Bogey'
  ELSE 'Triple Bogey or Worse'
END
```

In this example, the WHEN statements must reflect a strict equality. For example, a WHEN condition of WHEN < 0 THEN 'Under Par' is illegal because comparison operators are not allowed.

CASE (If)

This form of the CASE statement evaluates each WHEN condition and if satisfied, assigns the value in the corresponding THEN expression.

If none of the WHEN conditions are satisfied, it assigns the default value specified in the ELSE expression. If no ELSE expression is specified, the system automatically adds an ELSE NULL.

Note: In a CASE statement, AND has precedence over OR.

Syntax

```
CASE
```

```

    WHEN request_condition1 THEN expr1
    {WHEN request_condition2 THEN expr2}
    {WHEN request_condition... THEN expr...}
    ELSE expr
END

```

Where:

CASE starts the CASE statement. Must be followed by one or more WHEN and THEN statements, an optional ELSE statement, and the END keyword.

WHEN specifies the condition to be satisfied.

THEN specifies the value to assign if the corresponding WHEN expression is satisfied.

ELSE specifies the value to assign if none of the WHEN conditions are satisfied. If omitted, ELSE NULL is assumed.

END ends the CASE statement.

Example

```

CASE
    WHEN score-par < 0 THEN 'Under Par'
    WHEN score-par = 0 THEN 'Par'
    WHEN score-par = 1 THEN 'Bogie'
    WHEN score-par = 2 THEN 'Double Bogey'
    ELSE 'Triple Bogey or Worse'
END

```

Unlike the Switch form of the CASE statement, the WHEN statements in the If form allow comparison operators. For example, a WHEN condition of WHEN < 0 THEN 'Under Par' is legal.

Expressing Literals

A literal is a nonnull value corresponding to a given data type. Literals are typically constant values, or in other words, they are values that are taken as they are. A literal value must comply with the data type that it represents.

SQL provides mechanisms for expressing literals in SQL statements. This following topics describe how to express each type of literal in SQL:

- [Character Literals](#)
- [Datetime Literals](#)
- [Numeric Literals](#)

Character Literals

A character literal represents a value of CHARACTER or VARCHAR data type. To express a character literal, enclose the character string in single quotes ('). The number of characters enclosed between the single quotes implies the length of the literal.

Examples

```

'Oracle BI Server'
'abc123'

```

Datetime Literals

The SQL 92 standard defines three kinds of 'typed' datetime literals, in the following formats:

```
DATE 'yyyy-mm-dd'  
TIME 'hh:mm:ss'  
TIMESTAMP 'yyyy-mm-dd hh:mm:ss'
```

To express a typed datetime literal, use the keywords `DATE`, `TIME`, or `TIMESTAMP` followed by a datetime string enclosed in single quotation marks, as in the preceding example. Two digits are required for all nonyear components even if the value is a single digit.

These formats are fixed and are not affected by the format specified in the `NQSCONFIG.INI` file for the parameters `DATE_DISPLAY_FORMAT`, `TIME_DISPLAY_FORMAT`, or `DATE_TIME_DISPLAY_FORMAT`.

Examples

```
DATE '2000-08-15'  
TIME '11:55:25'  
TIMESTAMP '1999-03-15 11:55:25'
```

Numeric Literals

A numeric literal represents a value of a numeric data type (such as `INTEGER`, `DECIMAL`, or `FLOAT`). To express a numeric literal, type the number as part of a SQL statement.

Do not surround numeric literals with single quotes. Doing so expresses the literal as a character literal.

Numeric literals include:

- [Integer Literals](#)
- [Decimal Literals](#)
- [Floating Point Literals](#)

Integer Literals To express an integer constant as a literal, specify the integer as part of a SQL statement (for example, in the `SELECT` list). Precede the integer with a plus sign (+) to indicate the integer is positive, or a minus sign (-) to indicate the integer is negative. Unsigned integers are assumed to be positive.

Examples

```
234  
+2  
567934
```

Decimal Literals To express a decimal literal, specify a decimal number. Precede the number with a plus sign (+) to indicate the number is positive, or a minus sign (-) to indicate the number is negative. Unsigned numbers are assumed to be positive.

Examples

```
1.223  
-22.456  
+33.456789
```

Floating Point Literals To express floating point numbers as literal constants, enter a decimal literal followed by the letter E (either uppercase or lowercase), followed by the plus sign (+) to indicate a positive exponent, or the minus sign (-) to indicate a negative exponent. No spaces are allowed between the integer, the letter E, and the sign of the exponent.

Examples

```
333.456E-
1.23e+
```

Calculated Members

A calculated member is a user-defined dimension member whose measure values are calculated at run time.

You define a calculated member within a dimension through a formula that references other members of the same dimension. If a dimension has multiple hierarchies, all members referenced in the formula must belong to one hierarchy.

Within a calculated member, the members do not have to be at the same level in the hierarchy. For example, in a Geography hierarchy, you can create a calculated member to enable you to add together measure values for the Country member France and the City member Budapest.

The three standard components of a calculated member are:

- The presentation hierarchy on which the calculated member is based (for example, "Geography")
- The name to identify the calculated member, and to distinguish it from other members in the dimension (for example, 'My Locations')
- The formula used to calculate the calculated member, containing one or more Member clauses (for example, Member ("Geography"."Country"."France') + Member ("Geography"."City"."Budapest'))

This section contains the following topics:

- [CALCULATEDMEMBER Syntax](#)
- [Rules for the CALCULATEDMEMBER Expression](#)
- [Using Solve Order to Control Formula Evaluation Sequence](#)
- [Examples of Calculated Members in Queries](#)

CALCULATEDMEMBER Syntax

```
CALCULATEDMEMBER(presentation_hierarchy, member_identifier, calculated_member_  
formula [, solve_order])
```

Where:

presentation_hierarchy identifies the fully qualified presentation hierarchy in the presentation layer on which the calculated member is based, as follows:

```
"subject_area"."presentation_table"."presentation_hierarchy"
```

Note: When qualifying presentation hierarchies and presentation hierarchy levels in both the `CALCULATEDMEMBER` expression and the Member clause within the `calculated_member_formula` parameter, the following rule applies:

- You must specify the qualification term (`"subject_area".`) if there are multiple presentation tables or presentation hierarchies with the same name in different subject areas, otherwise you can omit the term.
-
-

`member_identifier` is the string or numeric literal that identifies the calculated member. The type of literal depends on the data type of the dimension level-keys.

`calculated_member_formula` consists of one or more examples of a "member clause" connected by the standard arithmetic operators `+ - * / ()`. The syntax of the member clause depends on whether the presentation hierarchy is level-based or parent-child. See ["Syntax for the Member Clause in Level-Based Hierarchies"](#) and ["Syntax for the Member Clause in Parent-Child Hierarchies"](#) for details.

`solve_order` (*optional*) is a positive integer, used to determine the order of evaluation when there are calculated members from different dimensions in the same query. See ["Using Solve Order to Control Formula Evaluation Sequence"](#) for details.

Syntax for the Member Clause in Level-Based Hierarchies

`MEMBER`(*presentation_hierarchy_level*, *member_value*)

Where:

`presentation_hierarchy_level` identifies the fully qualified hierarchy level in the `presentation_hierarchy`, as follows:

`"subject_area"."presentation_table"."presentation_hierarchy"."presentation_level"`

`member_value` is the string or numeric literal that identifies the member in the `presentation_hierarchy_level`.

Syntax for the Member Clause in Parent-Child Hierarchies

`MEMBER`(*presentation_hierarchy*, *member_value*)

Where:

`presentation_hierarchy` identifies the fully qualified presentation hierarchy in the presentation layer on which the calculated member is based, as follows:

`"subject_area"."presentation_table"."presentation_hierarchy"`

`member_value` is the string or numeric literal that identifies the member in the `presentation_hierarchy`.

Rules for the CALCULATEDMEMBER Expression

The rules for calculated members relate to the `CALCULATEDMEMBER` expression itself and the use of the `CALCULATEDMEMBER` expression in queries.

- All level references in a given `CALCULATEDMEMBER` expression must belong to the same dimension hierarchy.
- `CALCULATEDMEMBER` expressions may only appear in the `SELECT` list of a query.

- Only one `CALCULATEDMEMBER` expression is allowed for each dimension for each `SELECT` list of a query block. However, `CALCULATEDMEMBER` expressions based on other dimensions may exist in the same query.
- You cannot include any other column from a dimension on which a calculated member is based in the following components of a query block:
 - `SELECT` list
 - `WHERE` clause
 - `HAVING` clause

However, you may reference columns from the calculated member dimension in subqueries.

- Columns from other dimensions may be referenced in the same query block, as long as there are no `CALCULATEDMEMBER` expressions on those dimensions.

Using Solve Order to Control Formula Evaluation Sequence

By default, when the `CALCULATEDMEMBER` expression does not contain a solve order, the calculated members are evaluated in the order in which they appear in the `SELECT` list.

When there are calculated members from different dimensions in the same query block, the order in which the Oracle BI Server evaluates the calculated members may be significant.

[Example C-1](#) illustrates how the wrong solve order can lead to incorrect results.

Example C-1 Using Solve Order

Assume you have the following account and time data:

Kwik Grains	Profit	Sales
2007 Q3	300	1000
2007 Q4	600	1500

You want to calculate the percentage profit ($\text{Profit} / \text{Sales} * 100$) for each time period and the totals for the two quarters.

If the solve order for your calculations is the following:

1. `'Profit%' = 'Profit'/'Sales' * 100`
2. `'2007 Second Half' = '2007 Q3' + '2007 Q4'`

then the percentage profit for '2007 Second Half' is calculated *incorrectly* by adding the 2007 Q3 profit to the 2007 Q4 profit:

- $(300/1000) + (600/1500) = 30\% + 40\% = 70\%$

with the following results:

Kwik Grains	Profit	Sales	Profit%
2007 Q3	300	1000	30
2007 Q4	600	1500	40
2007 Second Half	900	2500	70 (<i>incorrect result</i>)

If the solve order for your calculations is the following:

1. '2007 Second Half' = '2007 Q3' + '2007 Q4'
2. 'Profit%' = 'Profit'/'Sales' * 100

then the percentage profit for '2007 Second Half' is calculated correctly by adding the 2007 Q3 and 2007 Q4 profits and sales first, then dividing the total Profit by the total Sales:

- $(300+600) / (1000+1500) = 900/2500 = 36\%$

with the following results:

Kwik Grains	Profit	Sales	Profit%
2007 Q3	300	1000	30
2007 Q4	600	1500	40
2007 Second Half	900	2500	36 (<i>correct result</i>)

See [Example C-3, "Using Calculated Members from Different Dimensions"](#) for an example of a query that explicitly specifies the solve order.

Examples of Calculated Members in Queries

The examples in this section show the use of calculated members in queries, and the base data on which the calculations are performed.

Example C-2 Single Calculated Member Query

This example shows two queries, each with corresponding results.

The first query contains a calculated member.

```
SELECT CALCULATEDMEMBER(product."Product - Region", 'USA - LA - Tokyo',
    MEMBER(product."Product - Region"."Country", 'USA')
    - MEMBER(product."Product - Region"."Region", 'LA')
    - MEMBER(product."Product - Region"."City", 'Tokyo')
) MyRegion,
    sales.Revenue Revenue, sales.QtySold QtySold
FROM product, sales;
```

Result:

MYREGION	REVENUE	QTYSOLD
USA - LA - Tokyo	61959.00	3959

The second query verifies the results of the first query. It shows the base data on which the calculation in the first query is performed.

```
SELECT * from SupplierCity where Country in ('USA', 'Japan');
```

Result:

CITY	REGION	COUNTRY	REVENUE	QTYSOLD
Boston	MA	USA	28146.40	2084
Osaka		Japan	15678.30	1417
New Orleans	LA	USA	33351.95	1735
Ann Arbor	MI	USA	43569.00	1436
Tokyo		Japan	33533.20	1134
Bend	OR	USA	23776.80	1573

Example C-3 Using Calculated Members from Different Dimensions

The requirement in this example is to determine the percentage increase over time in Revenue and Quantity Sold for US and Canada combined.

To achieve the correct results, the solve order is significant. You must first add Revenue and Quantity Sold for the two countries across the time periods, then perform the percentage calculation. See ["Using Solve Order to Control Formula Evaluation Sequence"](#) for more information about solve order significance.

This example shows two queries, each with corresponding results.

The first query contains the calculated members from the two dimensions Product - Region and Time, with the "addition" formula calculated first, then the "percentage" formula.

```
SELECT CALCULATEDMEMBER(product."Product - Region", 'North America',
    MEMBER(product."Product - Region"."Country", 'USA')
    + MEMBER(product."Product - Region"."Country", 'Canada'), 1
    ) MyRegion,
    CALCULATEDMEMBER(day."Time", 'Percentage Increase',
    ( MEMBER(day."Time"."Year", 1996)
    - MEMBER(day."Time"."Year", 1995) ) * 100
    / MEMBER(day."Time"."Year", 1995), 2
    ) MyTime,
    sales.Revenue RevenuePC,
    sales.QtySold QtySoldPC
FROM product, sales, day;
```

Result:

MYREGION	MYTIME	REVENUEPC	QTYSOLDPC
North America	Percentage Increase	16	35

Note that in the preceding query, the sequence of the calculated members in the SELECT list is sufficient for correct results, even without the explicit solve orders. The solve orders are included for completeness.

The second query verifies the results of the first query. It shows the base data on which the calculations in the first query are performed.

```
SELECT CALCULATEDMEMBER(product."Product - Region", 'North America',
    MEMBER(product."Product - Region"."Country", 'USA')
    + MEMBER(product."Product - Region"."Country", 'Canada')
    ) MyRegion,
    year as Year, sales.Revenue Revenue, sales.QtySold QtySold
FROM product, sales, day;
```

Result:

MYREGION	YEAR	REVENUE	QTYSOLD
North America	1996	101702.75	4918
North America	1995	87265.10	3638
North America	1994	30776.00	1616

Variables

You can include and set variables in SQL statements. To do this, include the variable at the beginning of the SQL statement.

Syntax

```
SET VARIABLE variable_name = variable_value; SELECT_statement
```

If you are executing a query from the nqcmd utility, use a colon as a delimiter. Otherwise, you can use either a semicolon or a colon.

Examples

```
SET VARIABLE LOGLEVEL = 3; SELECT Products.Brand, Measures.Dollars FROM "Products"
```

```
SET VARIABLE DISABLE_CACHE_HIT=1, LOGLEVEL = 3, WEBLANGUAGE='en': SELECT  
Products.Brand, Measures.Dollars FROM "Products"
```

Aggregate, Running Aggregate, and Time Series Functions

This section contains information about aggregate functions, running aggregate functions, and time series functions:

- [Aggregate Functions](#)
- [Running Aggregate Functions](#)
- [Time Series Functions](#)

Aggregate Functions

Aggregate functions perform operations on multiple values to create summary results.

The aggregate functions cannot be used to form nested aggregation in expressions on logical columns that have a default aggregation rule defined in the Aggregation tab of the Logical Column dialog. To specify nested aggregation, you must define a column with a default aggregation rule and then request the aggregation of the column in a SQL statement.

Aggregate functions include:

- [AGGREGATE AT](#)
- [AVG](#)
- [AVGDISTINCT](#)
- [BOTTOMN](#)
- [COUNT](#)
- [COUNTDISTINCT](#)
- [COUNT\(*\)](#)
- [FIRST](#)
- [GROUPBYCOLUMN](#)
- [GROUPBYLEVEL](#)
- [LAST](#)
- [MAX](#)
- [MEDIAN](#)
- [MIN](#)
- [NTILE](#)
- [PERCENTILE](#)
- [RANK](#)

- [STDDEV](#)
- [STDDEV_POP](#)
- [SUM](#)
- [SUMDISTINCT](#)
- [TOPN](#)

AGGREGATE AT

This function aggregates columns based on the level or levels you specify. Using `AGGREGATE AT` guarantees that the aggregate for the measure always occurs at the levels specified after the keyword `AT`, regardless of the `WHERE` clause.

Syntax

```
AGGREGATE(expr AT level [, level1, levelN])
```

Where:

expr is any expression that references at least one measure column

level is the level at which you want to aggregate. You can optionally specify multiple levels.

You cannot specify a level from a dimension that contains levels that are being used as the measure level for the measure you specified in the first argument. For example, you cannot write the function as `AGGREGATE(yearly_sales AT month)` because "month" is from the same time dimension that is being used as the measure level for "yearly_sales."

Example

The following example shows the `AGGREGATE AT` function and example results:

```
SELECT month, year, AGGREGATE(sales AT Year)
FROM timeseriestesting
WHERE year = 1994 AND month = 12
```

Result:

Month	Year	AGGREGATE AT year
12	1994	7396

Row count: 1

Because the `AGGREGATE AT` operator is always executed before the predicates, it always returns the correct total for the time level specified after the keyword `AT`.

AVG

This function calculates the average (mean) value of an expression in a result set. It must take a numeric expression as its argument.

Note that the denominator of `AVG` is the number of rows aggregated. For this reason, it is usually a mistake to use `AVG(x)` in a calculation in Oracle Business Intelligence. Instead, write the expression manually so that you can control both the numerator and denominator (x/y).

Syntax

```
AVG(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

AVGDISTINCT

This function calculates the average (mean) of all distinct values of an expression. It must take a numeric expression as its argument.

Syntax

```
AVG(DISTINCT numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

BOTTOMN

This function ranks the lowest *n* values of the expression argument from 1 to *n*, 1 corresponding to the lowest numeric value. The `BOTTOMN` function operates on the values returned in the result set. A request can contain only one `BOTTOMN` expression.

Syntax

```
BOTTOMN(numExpr, integer)
```

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer. Represents the bottom number of rankings displayed in the result set, 1 being the lowest rank.

COUNT

This function calculates the number of rows having a nonnull value for the expression. The expression is typically a column name, in which case the number of rows with nonnull values for that column is returned.

Syntax:

```
COUNT(expr)
```

Where:

expr is any expression.

COUNTDISTINCT

This function adds distinct processing to the `COUNT` function.

Syntax

```
COUNT(DISTINCT expr)
```

Where:

expr is any expression.

COUNT(*)

This function counts the number of rows.

Syntax

```
COUNT(*)
```

Example

For example, if a table named Facts contained 200,000,000 rows, the sample request would return the results shown:

```
SELECT COUNT(*) FROM Facts
```

Result:

```
200000000
```

FIRST

This function selects the first returned value of the expression argument. For example, the `FIRST` function can calculate the value of the first day of the year.

The `FIRST` function is limited to defining dimension-specific aggregation rules in a repository. You cannot use it in SQL statements.

The `FIRST` function operates at the most detailed level specified in your explicitly defined dimension. For example, if you have a time dimension defined with hierarchy levels day, month, and year, the `FIRST` function returns the first day in each level.

You should not use the `FIRST` function as the first dimension-specific aggregate rule. It might cause queries to bring back large numbers of rows for processing in the Oracle BI Server, causing poor performance.

When a measure is based on dimensions, and data is dense, the Oracle BI Server optimizes the SQL statements sent to the database to improve performance. See ["Setting Default Levels of Aggregation for Measure Columns"](#) for more information about dense data.

Note that you cannot nest `PERIODROLLING`, `FIRST`, and `LAST` functions.

Syntax

```
FIRST(expr)
```

Where:

expr is any expression that references at least one measure column.

Example

```
FIRST(sales)
```

GROUPBYCOLUMN

For use in setting up aggregate navigation. It specifies the logical columns that define the level of the aggregate data existing in a physical aggregate table.

For example, if an aggregate table contains data grouped by store and by month, specify the following syntax in the content filter (General tab of Logical Source dialog):

```
GROUPBYCOLUMN(STORE, MONTH)
```

The `GROUPBYCOLUMN` function is only for use in configuring a repository. You cannot use it to form SQL statements.

GROUPBYLEVEL

For use in setting up aggregate navigation. It specifies the dimension levels that define the level of the aggregate data existing in a physical aggregate table.

For example, if an aggregate table contains data at the store and month levels, and if you have defined dimensions (Geography and Customers) containing these levels, specify the following syntax in the content filter (General tab of Logical Source dialog):

```
GROUPBYLEVEL (GEOGRAPHY.STORE, CUSTOMERS.MONTH)
```

The `GROUPBYLEVEL` function is only for use in configuring a repository. You cannot use it to form SQL statements.

LAST

This function selects the last returned value of the expression. For example, the `LAST` function can calculate the value of the last day of the year.

The `FIRST` function is limited to defining dimension-specific aggregation rules in a repository. You cannot use it in SQL statements.

The `LAST` function operates at the most detailed level specified in your explicitly defined dimension. For example, if you have a time dimension defined with hierarchy levels day, month, and year, the `LAST` function returns the last day in each level.

You should not use the `LAST` function as the first dimension-specific aggregate rule. It might cause queries to bring back large numbers of rows for processing in the Oracle BI Server, causing poor performance.

When a measure is based on dimensions, and data is dense, the Oracle BI Server optimizes the SQL statements sent to the database to improve performance. See "[Setting Default Levels of Aggregation for Measure Columns](#)" for more information about dense data.

Note that you cannot nest `PERIODROLLING`, `FIRST`, and `LAST` functions.

Syntax

```
LAST (expr)
```

Where:

expr is any expression that references at least one measure column.

Example

```
LAST(sales)
```

MAX

This function calculates the maximum value (highest numeric value) of the rows satisfying the numeric expression argument.

Syntax

```
MAX (numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The `MAX` function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

MEDIAN

This function calculates the median (middle) value of the rows satisfying the numeric expression argument. When there are an even number of rows, the median is the mean of the two middle rows. This function always returns a double.

Syntax

```
MEDIAN(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The MEDIAN function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

MIN

This function calculates the minimum value (lowest numeric value) of the rows satisfying the numeric expression argument.

Syntax

```
MIN(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The MIN function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

NTILE

This function determines the rank of a value in terms of a user-specified range. It returns integers to represent any range of ranks. In other words, the resulting sorted data set is broken into several tiles where there are roughly an equal number of values in each tile.

NTile with *numTiles* = 100 returns what is commonly called the "percentile" (with numbers ranging from 1 to 100, with 100 representing the high end of the sort). This value is different from the results of the Oracle BI PERCENTILE function, which conforms to what is called "percent rank" in SQL 92 and returns values from 0 to 1.

Syntax

```
NTILE(numExpr, numTiles)
```

Where:

numExpr is any expression that evaluates to a numeric value.

numTiles is a positive, nonnull integer that represents the number of tiles.

If the *numExpr* argument is not null, the function returns an integer that represents a rank within the requested range.

PERCENTILE

This function calculates a percent rank for each value satisfying the numeric expression argument. The percentile rank ranges are from 0 (1st percentile) to 1 (100th percentile), inclusive.

The percentile is calculated based on the values in the result set.

Syntax

```
PERCENTILE(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The PERCENTILE function resets its values for each group in the query according to specific rules. See ["Display Function Reset Behavior"](#) for more information.

RANK

This function calculates the rank for each value satisfying the numeric expression argument. The highest number is assigned a rank of 1, and each successive rank is assigned the next consecutive integer (2, 3, 4,...). If certain values are equal, they are assigned the same rank (for example, 1, 1, 1, 4, 5, 5, 7...).

The rank is calculated based on the values in the result set.

Syntax

```
RANK(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The RANK function resets its values for each group in the query according to specific rules. See ["Display Function Reset Behavior"](#) for more information.

STDDEV

This function returns the standard deviation for a set of values. The return type is always a double. STDEV_SAMP is a synonym for STDDEV.

Syntax

```
STDDEV([ALL | DISTINCT] numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

If ALL is specified, the standard deviation is calculated for all data in the set.

If DISTINCT is specified, all duplicates are ignored in the calculation.

If nothing is specified (the default), all data is considered.

The STDDEV function resets its values for each group in the query according to specific rules. See ["Display Function Reset Behavior"](#) for more information.

STDDEV_POP

This function returns the standard deviation for a set of values using the computational formula for population variance and standard deviation.

Syntax

```
STDDEV_POP([ALL | DISTINCT] numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

If ALL is specified, the standard deviation is calculated for all data in the set.

If DISTINCT is specified, all duplicates are ignored in the calculation.

If nothing is specified (the default), all data is considered.

SUM

This function calculates the sum obtained by adding up all values satisfying the numeric expression argument.

Syntax

```
SUM(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The SUM function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

SUMDISTINCT

This function calculates the sum obtained by adding all of the distinct values satisfying the numeric expression argument.

Syntax

```
SUM(DISTINCT numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

TOPN

This function ranks the highest *n* values of the expression argument from 1 to *n*, 1 corresponding to the highest numeric value. The TOPN function operates on the values returned in the result set. A request can contain only one TOPN expression.

Syntax

```
TOPN(numExpr, integer)
```

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer. Represents the top number of rankings displayed in the result set, 1 being the highest rank.

The TOPN function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

Running Aggregate Functions

Running aggregate functions are similar to functional aggregates in that they take a set of records as input, but instead of outputting the single aggregate for the entire set of records, they output the aggregate based on records encountered so far.

This section describes the running aggregate functions supported by the Oracle BI Server. Functions include:

- [MAVG](#)
- [MSUM](#)
- [RSUM](#)
- [RCOUNT](#)
- [RMAX](#)
- [RMIN](#)

MAVG

This function calculates a moving average (mean) for the last n rows of data in the result set, inclusive of the current row.

The average for the first row is equal to the numeric expression for the first row. The average for the second row is calculated by taking the average of the first two rows of data. The average for the third row is calculated by taking the average of the first three rows of data, and so on until you reach the n th row, where the average is calculated based on the last n rows of data.

Syntax

`MAVG(numExpr, integer)`

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer. Represents the average of the last n rows of data.

The `MAVG` function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

MSUM

This function calculates a moving sum for the last n rows of data, inclusive of the current row.

The sum for the first row is equal to the numeric expression for the first row. The sum for the second row is calculated by taking the sum of the first two rows of data. The sum for the third row is calculated by taking the sum of the first three rows of data, and so on. When the n th row is reached, the sum is calculated based on the last n rows of data.

Syntax

`MSUM(numExpr, integer)`

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer. Represents the average of the last n rows of data.

The `MSUM` function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

Example

This example shows a query that uses the MSUM function, along with example query results.

```
select month, revenue, MSUM(revenue, 3) as 3_MO_SUM from sales_subject_area
```

Result:

MONTH	REVENUE	3_MO_SUM
JAN	100.00	100.00
FEB	200.00	300.00
MAR	100.00	400.00
APRIL	100.00	400.00
MAY	300.00	500.00
JUNE	400.00	800.00
JULY	500.00	1200.00
AUG	500.00	1400.00
SEPT	500.00	1500.00
OCT	300.00	1300.00
NOV	200.00	1000.00
DEC	100.00	600.00

RSUM

This function calculates a running sum based on records encountered so far. The sum for the first row is equal to the numeric expression for the first row. The sum for the second row is calculated by taking the sum of the first two rows of data. The sum for the third row is calculated by taking the sum of the first three rows of data, and so on.

Syntax

```
RSUM(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

The RSUM function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

Example

This example shows a query that uses the RSUM function, along with example query results.

```
SELECT month, revenue, RSUM(revenue) as RUNNING_SUM from sales_subject_area
```

Result:

MONTH	REVENUE	RUNNING_SUM
JAN	100.00	100.00
FEB	200.00	300.00
MAR	100.00	400.00
APRIL	100.00	500.00
MAY	300.00	800.00
JUNE	400.00	1200.00
JULY	500.00	1700.00
AUG	500.00	2200.00
SEPT	500.00	2700.00
OCT	300.00	3000.00
NOV	200.00	3200.00
DEC	100.00	3300.00

RCOUNT

This function takes a set of records as input and counts the number of records encountered so far.

Syntax

`RCOUNT(expr)`

Where:

expr is an expression of any data type.

The RCOUNT function resets its values for each group in the query according to specific rules. See ["Display Function Reset Behavior"](#) for more information.

Example

This example shows a query that uses the RCOUNT function, along with example query results.

```
select month, profit, RCOUNT(profit) from sales_subject_area where profit > 200
```

Result:

MONTH	PROFIT	RCOUNT(profit)
MAY	300.00	2
JUNE	400.00	3
JULY	500.00	4
AUG	500.00	5
SEPT	500.00	6
OCT	300.00	7

RMAX

This function takes a set of records as input and shows the maximum value based on records encountered so far. The specified data type must be one that can be ordered.

Syntax

`RMAX(expr)`

Where:

expr is an expression of any data type. The data type must be one that has an associated sort order.

The RMAX function resets its values for each group in the query according to specific rules. See ["Display Function Reset Behavior"](#) for more information.

Example

This example shows a query that uses the RMAX function, along with example query results.

```
SELECT month, profit, RMAX(profit) from sales_subject_area
```

Result:

MONTH	PROFIT	RMAX(profit)
JAN	100.00	100.00
FEB	200.00	200.00
MAR	100.00	200.00
APRIL	100.00	200.00

MAY	300.00	300.00
JUNE	400.00	400.00
JULY	500.00	500.00
AUG	500.00	500.00
SEPT	500.00	500.00
OCT	300.00	500.00
NOV	200.00	500.00
DEC	100.00	500.00

RMIN

This function takes a set of records as input and shows the minimum value based on records encountered so far. The specified data type must be one that can be ordered.

Syntax

`RMIN(expr)`

Where:

expr is an expression of any data type. The data type must be one that has an associated sort order.

The RMIN function resets its values for each group in the query according to specific rules. See "[Display Function Reset Behavior](#)" for more information.

Example

This example shows a query that uses the RMIN function, along with example query results.

```
select month, profit, RMIN(profit) from sales_subject_area
```

Result:

MONTH	PROFIT	RMIN(profit)
JAN	400.00	400.00
FEB	200.00	200.00
MAR	100.00	100.00
APRIL	100.00	100.00
MAY	300.00	100.00
JUNE	400.00	100.00
JULY	500.00	100.00
AUG	500.00	100.00
SEPT	500.00	100.00
OCT	300.00	100.00
NOV	200.00	100.00
DEC	100.00	100.00

Time Series Functions

Time series functions operate on time-oriented dimensions. The time series functions calculate `AGO`, `TODATE`, and `PERIODROLLING` functions based on user supplied calendar tables, not on standard SQL date manipulation functions.

These functions let you use Expression Builder to call a logical function to perform time series calculations instead of aliasing physical tables and modeling logically.

To use time series functions on a particular dimension, you must designate the dimension as a Time dimension and set one or more keys at one or more levels as chronological keys. See "[Modeling Time Series Data](#)" for more information.

Functions include:

- [AGO](#)
- [PERIODROLLING](#)
- [TODATE](#)

AGO

This function is a time series aggregation function that calculates the aggregated value from the current time back to a specified time period. For example, `AGO` can produce sales for every month of the current quarter and the corresponding quarter-ago sales.

Time series functions operate on members of time dimensions which are at or below the level of the function. Because of this, one or more columns that uniquely identify members at or below the given level must be projected in the query. Alternatively, you can apply a filter to the query that specifies a single member at or below the given level. See "[Determining the Level Used by the AGO Function](#)" for more information about the level of the function.

If unsupported metrics are requested, NULL values are returned and a warning entry is written to the `nquery.log` file when the logging level equals three or above.

Multiple `AGO` functions can be nested if all the `AGO` functions have the same level argument. You can nest exactly one `TODATE` and multiple `AGO` functions if they each have the same level argument.

Syntax

```
AGO(expr, [time_level], offset)
```

Where:

expr is an expression that references at least one measure column.

time_level is an optional argument that specifies the type of time period, such as quarter, month, or year.

offset is an integer literal that represents the time shift amount.

Example

The following example returns last year's sales:

```
SELECT Year_ID, AGO(sales, year, 1)
```

Determining the Level Used by the AGO Function The unit of time (offset) used in the `AGO` function is called the *level* of the function. This value is determined by the measure level of the measures in its first argument, the `AGO` level (optionally specified within the function), and the query level of the query to which the function belongs.

- The measure level for the measure can be set in the Administration Tool. If a measure level has been set for the measure used in the function, the measure level is used as the level of the function. The measure level is also called the *storage grain* of the function.
- The `AGO` level can be optionally specified as the second argument of the function. If a measure level has not been set in the Administration Tool, but an `AGO` level has been specified, the `AGO` level is used as the level of the function. The `AGO` level is also called the *time series grain* of the function.

- If a measure level has not been set in the Administration Tool, and if no AGO level has been set explicitly in the function, the query level is used as the level of the function. The query level is also called the *query grain* of the function.

PERIODROLLING

This function computes the aggregate of a measure over the period starting x units of time and ending y units of time from the current time. For example, you can use PERIODROLLING to compute sales for a period that starts at a certain quarter before and ends at a certain quarter after the current quarter.

Time series functions operate on members of time dimensions which are at or below the level of the function. Because of this, one or more columns that uniquely identify members at or below the given level must be projected in the query. Alternatively, you can apply a filter to the query that specifies a single member at or below the given level. See "[Determining the Level Used by the PERIODROLLING Function](#)" for more information about the level of the function.

You cannot nest AGO and TODATE functions within a PERIODROLLING function. Also, you cannot nest PERIODROLLING, FIRST, and LAST functions.

If you embed other aggregate functions (like RANK, TOPN, PERCENTILE, FILTER, or RSUM) inside PERIODROLLING, the PERIODROLLING function is pushed inward. For example, PERIODROLLING(TOPN(*measure*)) is executed as TOPN(PERIODROLLING(*measure*)).

Syntax

```
PERIODROLLING(measure, x, y [, hierarchy])
```

Where:

measure is the name of a measure column.

x is an integer that specifies the offset from the current time. Precede the integer with a minus sign (-) to indicate an offset into the past.

y specifies the number of time units over which the function will compute. To specify the current time, enter 0.

hierarchy is an optional argument that specifies the name of a hierarchy in a time dimension, such as *yr*, *mon*, *day*, that you want to use to compute the time window. This option is useful when there are multiple hierarchies in a time dimension, or when you want to distinguish between multiple time dimensions.

If you want to roll back or forward the maximum possible amount, use the keyword UNBOUND. For example, the function PERIODROLLING(*measure*, -UNBOUND, 0) sums over the period starting from the beginning of time until now.

You can combine PERIODROLLING and AGGREGATE AT functions to specify the level of the PERIODROLLING function explicitly. For example, if the query level is day but you want to find the sum of the previous and current months, use the following:

```
SELECT year, month, day, PERIODROLLING(AGGREGATE(sales AT month), -1)
```

Examples

```
SELECT Month_ID, PERIODROLLING(monthly_sales, -1, 1)
```

```
SELECT Month_ID, PERIODROLLING(monthly_sales, -UNBOUND, 2)
```

```
SELECT Month_ID, PERIODROLLING(monthly_sales, -UNBOUND, UNBOUND)
```

Determining the Level Used by the PERIODROLLING Function The unit of time (offset) used in the PERIODROLLING function is called the *level* of the function. This value is determined by the measure level of the measures in its first argument and the query level of the query to which the function belongs. The measure level for the measure can be set in the Administration Tool. If a measure level has been set for the measure used in the function, the measure level is used as the level of the function. The measure level is also called the *storage grain* of the function.

If a measure level has not been set in the Administration Tool, then the query level is used. The query level is also called the *query grain* of the function. In the following example, the query level is month, and the PERIODROLLING function computes the sum of the last, current, and next month for each city for the months of March and April:

```
SELECT year, month, country, city, PERIODROLLING(sales, -1, 1)
WHERE month in ('Mar', 'Apr') AND city = 'New York'
```

When there are multiple hierarchies in the time dimension, you must specify the *hierarchy* argument in the PERIODROLLING function. For example:

```
SELECT year, fiscal_year, month, PERIODROLLING(sales, -1, 1, "fiscal_time_
hierarchy")
```

In this example, the level of the PERIODROLLING function is `fiscal_year`.

TODATE

This function is a time series aggregation function that aggregates a measure from the beginning of a specified time period to the currently displayed time. For example, this function can calculate Year to Date sales.

Time series functions operate on members of time dimensions which are at or below the level specified in the function. Because of this, one or more columns that uniquely identify members at or below the given level must be projected in the query. Alternatively, you can apply a filter to the query that specifies a single member at or below the given level.

If unsupported metrics are requested, NULL values are returned and a warning entry is written to the `nquery.log` file when the logging level equals three or above.

A TODATE function may not be nested within another TODATE function. You can nest exactly one TODATE and multiple AGO functions if they each have the same level argument.

TODATE is different from the TO_DATE SQL function supported by some databases. Do not use TO_DATE to change to a DATE data type. Instead, use the CAST function. See "[CAST](#)" for more information.

Syntax

```
TODATE(expr, time_level)
```

Where:

expr is an expression that references at least one measure column.

time_level is the type of time period, such as quarter, month, or year.

Example

The following example returns the year-to-month sales:

```
SELECT Year_ID, Month_ID, TODATE(sales, year)
```

String Functions

String functions perform various character manipulations, and they operate on character strings. Functions include:

- ASCII
- BIT_LENGTH
- CHAR
- CHAR_LENGTH
- CONCAT
- INSERT
- LEFT
- LENGTH
- LOCATE
- LOCATEN
- LOWER
- OCTET_LENGTH
- POSITION
- REPEAT
- REPLACE
- RIGHT
- SPACE
- SUBSTRING
- TRIMBOTH
- TRIMLEADING
- TRIMTRAILING
- UPPER

ASCII

This function converts a single character string to its corresponding ASCII code, between 0 and 255. If the character expression evaluates to multiple characters, the ASCII code corresponding to the first character in the expression is returned.

Syntax

```
ASCII(strExpr)
```

Where:

strExpr is any expression that evaluates to a character string.

BIT_LENGTH

This function returns the length, in bits, of a specified string. Each Unicode character is 2 bytes in length (equal to 16 bits).

Syntax

```
BIT_LENGTH(strExpr)
```

Where:

strExpr is any expression that evaluates to character string.

CHAR

This function converts a numeric value between 0 and 255 to the character value corresponding to the ASCII code.

Syntax

```
CHAR(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value between 0 and 255.

CHAR_LENGTH

This function returns the length, in number of characters, of a specified string. Leading and trailing blanks are not counted in the length of the string.

Syntax

```
CHAR_LENGTH(strExpr)
```

Where:

strExpr is any expression that evaluates to a character string.

CONCAT

There are two forms of this function. The first form concatenates two character strings. The second form uses the character string concatenation character to concatenate more than two character strings.

Syntax for Form 1 (To Concatenate Two Strings)

```
CONCAT(strExpr1, strExpr2)
```

Where:

strExprs are expressions that evaluate to character strings, separated by commas.

Example

This example request returns the results shown.

```
SELECT DISTINCT CONCAT('abc', 'def') FROM employee  
CONCAT('abc', 'def')
```

Result:

```
abcdef
```

Syntax for Form 2 (To Concatenate More Than Two Strings)

```
CONCAT(strExpr1, strExpr2 || strExpr3)
```

Where:

strExprs are expressions that evaluate to character strings, separated by commas and the character string concatenation operator || (double vertical bars). First, *strExpr2* is concatenated with *strExpr3* to produce an intermediate string, then both *strExpr1* and the intermediate string are concatenated by the CONCAT function to produce the final string.

Example

This example request returns the results shown.

```
SELECT DISTINCT CONCAT('abc','def' || 'ghi') FROM employee
```

Result:

```
abcdefghijkl
```

INSERT

This function inserts a specified character string into a specified location in another character string.

Syntax

```
INSERT(strExpr1, integer1, integer2, strExpr2)
```

Where:

strExpr1 is any expression that evaluates to a character string. Identifies the target character string.

integer1 is any positive integer that represents the number of characters from the beginning of the target string where the second string is to be inserted.

integer2 is any positive integer that represents the number of characters in the target string to be replaced by the second string.

strExpr2 is any expression that evaluates to a character string. Identifies the character string to be inserted into the target string.

Example

In the first string, starting at the second position (occupied by the number 2), three characters (the numbers 2, 3, and 4) are replaced by the string abcd.

```
SELECT INSERT('123456', 2, 3, 'abcd') FROM table
```

Result:

```
1abcd56
1abcd56
...
```

LEFT

Returns a specified number of characters from the left of a string.

Syntax

```
LEFT(strExpr, integer)
```

Where:

strExpr is any expression that evaluates to a character string.

integer is any positive integer that represents the number of characters from the left of the string to return.

Example

This example returns the three leftmost characters from the character string 123456:

```
SELECT LEFT('123456', 3) FROM table
```

Result:

```
123  
123  
...
```

LENGTH

This function returns the length, in number of characters, of a specified string. The length is returned excluding any trailing blank characters.

Syntax

```
LENGTH(strExpr)
```

Where:

strExpr is any expression that evaluates to a character string.

LOCATE

This function returns the numeric position of a character string in another character string. If the character string is not found in the string being searched, the function returns a value of 0.

If you want to specify a starting position to begin the search, use the `LOCATEN` function instead. See "[LOCATEN](#)" for details.

Syntax

```
LOCATE(strExpr1, strExpr2)
```

Where:

strExpr1 is any expression that evaluates to a character string. Identifies the string for which to search.

strExpr2 is any expression that evaluates to a character string. Identifies the string to be searched.

Examples

This example returns 4 as the numeric position of the letter d in the character string abcdef:

```
Locate('d', 'abcdef')
```

This example returns 0, because the letter g is not found within the string being searched.

```
Locate('g', 'abcdef')
```

LOCATEN

This function returns the numeric position of a character string in another character string. LOCATEN is identical to the LOCATE function, except that the search begins at the position specified by an integer argument. If the character string is not found in the string being searched, the function returns a value of 0. The numeric position to return is determined by counting the first character in the string as occupying position 1, regardless of the value of the integer argument.

Syntax

```
LOCATEN(strExpr1, strExpr2, integer)
```

Where:

strExpr1 is any expression that evaluates to a character string. Identifies the string for which to search.

strExpr2 is any expression that evaluates to a character string. Identifies the string to be searched.

integer is any positive (nonzero) integer that represents the starting position to begin to look for the character string.

Examples

This example returns 4 as the numeric position of the letter d in the character string abcdef. The search begins with the letter c, the third character in the string. The numeric position to return is determined by counting the letter 'a' as occupying position 1.

```
LOCATEN('d' 'abcdef', 3)
```

This example returns 0, because the letter b occurs in the string before the starting position to begin the search.

```
LOCATEN('b' 'abcdef', 3)
```

LOWER

This function converts a character string to lowercase.

Syntax

```
LOWER(strExpr)
```

Where:

strExpr is any expression that evaluates to a character string.

OCTET_LENGTH

This function returns the number of bits, in base 8 units (number of bytes), of a specified string.

Syntax

```
OCTET_LENGTH(strExpr)
```

Where:

strExpr is any expression that evaluates to a character string.

POSITION

This function returns the numeric position of *strExpr1* in a character expression. If *strExpr1* is not found, the function returns 0. See also "LOCATE" and "LOCATEN" for related information.

Syntax

```
POSITION(strExpr1 IN strExpr2)
```

Where:

strExpr1 is any expression that evaluates to a character string. Identifies the string to search for in the target string.

strExpr2 is any expression that evaluates to a character string. Identifies the target string to be searched.

Examples

This example returns 4 as the position of the letter d in the character string abcdef:

```
POSITION('d', 'abcdef')
```

This example returns 0 as the position of the number 9 in the character string 123456, because the number 9 is not found.

```
POSITION('9', '123456')
```

REPEAT

This function repeats a specified expression *n* times.

Syntax

```
REPEAT(strExpr, integer)
```

Where:

strExpr is any expression that evaluates to a character string.

integer is any positive integer that represents the number of times to repeat the character string.

Example

This example repeats abc four times:

```
REPEAT('abc', 4)
```

REPLACE

This function replaces one or more characters from a specified character expression with one or more other characters.

Syntax


```
REPLACE(strExpr1, strExpr2, strExpr3)
```

Where:

strExpr1 is any expression that evaluates to a character string. This is the string in which characters are to be replaced.

strExpr2 is any expression that evaluates to a character string. This second string identifies the characters from the first string that are to be replaced.

strExpr3 is any expression that evaluates to a character string. This third string specifies the characters to substitute into the first string.

Example

In the character string `abcd1234`, the characters `123` are replaced by the character string `zz`:

```
Replace('abcd1234', '123', 'zz')
```

Result:

```
abcdzz4
```

RIGHT

This function returns a specified number of characters from the right of a string.

Syntax

```
RIGHT(strExpr, integer)
```

Where:

strExpr is any expression that evaluates to a character string.

integer is any positive integer that represents the number of characters from the right of the string to return.

Example

This example returns the three rightmost characters from the character string `123456`:

```
SELECT right('123456', 3) FROM table
```

Result:

```
456
```

SPACE

This function inserts blank spaces.

Syntax

```
SPACE(integer)
```

Where:

integer is any positive integer that indicates the number of spaces to insert.

SUBSTRING

This function creates a new string starting from a fixed number of characters into the original string.

Syntax

```
SUBSTRING(strExpr FROM starting_position)
```

Where:

strExpr is any expression that evaluates to a character string.

starting_position is any positive integer that represents the number of characters from the start of the left side of the string where the result is to begin.

TRIMBOTH

This function strips specified leading and trailing characters from a character string.

Syntax

```
TRIM(BOTH character FROM strExpr)
```

Where:

character is any single character. If you omit this specification (and the required single quotes), a blank character is used as the default.

strExpr is any expression that evaluates to a character string.

TRIMLEADING

This function strips specified leading characters from a character string.

Syntax

```
TRIM(LEADING character FROM strExpr)
```

Where:

character is any single character. If you omit this specification (and the required single quotes), a blank character is used as the default.

strExpr is any expression that evaluates to a character string.

TRIMTRAILING

This function strips specified trailing characters from a character string.

Syntax

```
TRIM(TRAILING character FROM strExpr)
```

Where:

character is any single character. If you omit this specification (and the required single quotes), a blank character is used as the default.

strExpr is any expression that evaluates to a character string.

UPPER

This function converts a character string to uppercase.

Syntax

`UPPER(strExpr)`

Where:

strExpr is any expression that evaluates to a character string.

Math Functions

The math functions perform mathematical operations. Functions include:

- [ABS](#)
- [ACOS](#)
- [ASIN](#)
- [ATAN](#)
- [ATAN2](#)
- [CEILING](#)
- [COS](#)
- [COT](#)
- [DEGREES](#)
- [EXP](#)
- [EXTRACTBIT](#)
- [FLOOR](#)
- [LOG](#)
- [LOG10](#)
- [MOD](#)
- [PI](#)
- [POWER](#)
- [RADIANs](#)
- [RAND](#)
- [RANDFROMSEED](#)
- [ROUND](#)
- [SIGN](#)
- [SIN](#)
- [SQRT](#)
- [TAN](#)
- [TRUNCATE](#)

ABS

This function calculates the absolute value of a numeric expression.

Syntax

ABS (*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

ACOS

This function calculates the arc cosine of a numeric expression.

Syntax

ACOS (*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

ASIN

This function calculates the arc sine of a numeric expression.

Syntax

ASIN (*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

ATAN

This function calculates the arc tangent of a numeric expression.

Syntax

ATAN (*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

ATAN2

This function calculates the arc tangent of y/x , where y is the first numeric expression and x is the second numeric expression.

Syntax

ATAN2 (*numExpr1*, *numExpr2*)

Where:

numExpr is any expression that evaluates to a numeric value.

CEILING

This function rounds a noninteger numeric expression to the next highest integer. If the numeric expression evaluates to an integer, the `CEILING` function returns that integer.

Syntax

`CEILING (numExpr)`

Where:

numExpr is any expression that evaluates to a numeric value.

COS

This function calculates the cosine of a numeric expression.

Syntax

`COS (numExpr)`

Where:

numExpr is any expression that evaluates to a numeric value.

COT

This function calculates the cotangent of a numeric expression.

Syntax

`COT (numExpr)`

Where:

numExpr is any expression that evaluates to a numeric value.

DEGREES

This function converts an expression from radians to degrees.

Syntax

`DEGREES (numExpr)`

Where:

numExpr is any expression that evaluates to a numeric value.

EXP

This function sends the value to the power specified.

Syntax

`EXP (numExpr)`

Where:

numExpr is any expression that evaluates to a numeric value.

EXTRACTBIT

This function retrieves a bit at a particular position in an integer. It returns an integer of either 0 or 1 corresponding to the position of the bit. The primary use case for this function is to extract 'cell status' in the Hyperion Financial Management cube source. The `EXTRACTBIT` function cannot be pushed into any database, and is always internally executed (in the Oracle BI Server).

Syntax

```
Int ExtractBit(Arg1, Arg2)
```

Where:

Arg1 is an expression of the following types: `INT`, `SMALLINT`, `UNIT`, `SMALLUNIT`, `TINYINT`, `TINYUNIT`. If *Arg1* is of double type, it is necessary to cast the column to an `INT` first.

Arg2 is an expression of type integer. The value should range from 1 to *length_of_Arg1*. 1 retrieves the Least Significant Bit. If the *Arg2* is beyond the length of the integer, then 0 is returned. An error message is triggered when the *Arg2* is less than 1.

FLOOR

This function rounds a noninteger numeric expression to the next lowest integer. If the numeric expression evaluates to an integer, the `FLOOR` function returns that integer.

Syntax

```
FLOOR(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

LOG

This function calculates the natural logarithm of an expression.

Syntax

```
LOG(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

LOG10

This function calculates the base 10 logarithm of an expression.

Syntax

```
LOG10(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

MOD

This function divides the first numeric expression by the second numeric expression and returns the remainder portion of the quotient.

Syntax

```
MOD(numExpr1, numExpr2)
```

Where:

numExpr is any expression that evaluates to a numeric value.

Examples

This example request returns a value of 0:

```
MOD(9, 3)
```

This example request returns a value of 1:

```
MOD(10, 3)
```

PI

This function returns the constant value of pi (the circumference of a circle divided by its diameter).

Syntax

```
PI()
```

POWER

This function takes the first numeric expression and raises it to the power specified in the second numeric expression.

Syntax

```
POWER(numExpr1, numExpr2)
```

Where:

numExpr1 is any expression that evaluates to a numeric value.

RADIANS

This function converts an expression from degrees to radians.

Syntax

```
RADIANS(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

RAND

Returns a pseudo-random number between 0 and 1.

Syntax

RAND()

RANDFROMSEED

Returns a pseudo-random number based on a seed value. For a given seed value, the same set of random numbers are generated.

Syntax

RAND(*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

ROUND

This function rounds a numeric expression to *n* digits of precision.

Syntax

ROUND(*numExpr*, *integer*)

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer that represents the number of digits of precision.

Example

This example returns 2.17 as the result.

```
ROUND(2.166000, 2)
```

SIGN

This function returns the following:

- A value of 1 if the numeric expression argument evaluates to a positive number.
- A value of -1 if the numeric expression argument evaluates to a negative number.
- 0 (zero) if the numeric expression argument evaluates to zero.

Syntax

SIGN(*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

SIN

This function calculates the sine of a numeric expression.

Syntax

SIN(*numExpr*)

Where:

numExpr is any expression that evaluates to a numeric value.

SQRT

This function calculates the square root of the numeric expression argument. The numeric expression must evaluate to a nonnegative number.

Syntax

```
SQRT(numExpr)
```

Where:

numExpr is any expression that evaluates to a nonnegative numeric value.

TAN

This function calculates the tangent of a numeric expression.

Syntax

```
TAN(numExpr)
```

Where:

numExpr is any expression that evaluates to a numeric value.

TRUNCATE

This function truncates a decimal number to return a specified number of places from the decimal point.

Syntax

```
TRUNCATE(numExpr, integer)
```

Where:

numExpr is any expression that evaluates to a numeric value.

integer is any positive integer that represents the number of characters to the right of the decimal place to return.

Examples

This example returns 45.12:

```
TRUNCATE(45.12345, 2)
```

This example returns 25.12:

```
TRUNCATE(25.126, 2)
```

Calendar Date/Time Functions

The calendar date/time functions manipulate data of the data types DATE and DATETIME based on a calendar year. You must select these functions with another column; they cannot be selected alone. Functions include:

- [CURRENT_DATE](#)
- [CURRENT_TIME](#)

- CURRENT_TIMESTAMP
- DAY_OF_QUARTER
- DAYNAME
- DAYOFMONTH
- DAYOFWEEK
- DAYOFYEAR
- HOUR
- MINUTE
- MONTH
- MONTH_OF_QUARTER
- MONTHNAME
- NOW
- QUARTER_OF_YEAR
- SECOND
- TIMESTAMPADD
- TIMESTAMPDIF
- WEEK_OF_QUARTER
- WEEK_OF_YEAR
- YEAR

CURRENT_DATE

This function returns the current date. The date is determined by the system in which the Oracle BI Server is running.

Syntax

```
CURRENT_DATE
```

CURRENT_TIME

This function returns the current time. The time is determined by the system in which the Oracle BI Server is running.

Syntax

```
CURRENT_TIME(integer)
```

Where:

integer is any integer representing the number of digits of precision with which to display the fractional second. The argument is optional; the function returns the default precision when no argument is specified.

CURRENT_TIMESTAMP

This function returns the current date/timestamp. The timestamp is determined by the system in which the Oracle BI Server is running.

Syntax

```
CURRENT_TIMESTAMP(integer)
```

Where:

integer is any integer representing the number of digits of precision with which to display the fractional second. The argument is optional; the function returns the default precision when no argument is specified.

DAY_OF_QUARTER

This function returns a number (between 1 and 92) corresponding to the day of the quarter for the specified date.

Syntax

```
DAY_OF_QUARTER(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

DAYNAME

This function returns the name of the day of the week for a specified date.

Syntax

```
DAYNAME(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

DAYOFMONTH

This function returns the number corresponding to the day of the month for a specified date.

Syntax

```
DAYOFMONTH(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

DAYOFWEEK

This function returns a number between 1 and 7 corresponding to the day of the week, Sunday through Saturday, for a specified date. For example, the number 1 corresponds to Sunday, and the number 7 corresponds to Saturday.

Syntax

```
DAYOFWEEK(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

DAYOFYEAR

This function returns the number (between 1 and 366) corresponding to the day of the year for a specified date.

Syntax

`DAYOFYEAR(dateExpr)`

Where:

dateExpr is any expression that evaluates to a date.

HOUR

This function returns a number (between 0 and 23) corresponding to the hour for a specified time. For example, 0 corresponds to 12 a.m. and 23 corresponds to 11 p.m.

Syntax

`HOUR(timeExpr)`

Where:

timeExpr is any expression that evaluates to a time.

MINUTE

This function returns a number (between 0 and 59) corresponding to the minute for a specified time.

Syntax

`MINUTE(timeExpr)`

Where:

timeExpr is any expression that evaluates to a time.

MONTH

This function returns the number (between 1 and 12) corresponding to the month for a specified date.

Syntax

`MONTH(dateExpr)`

Where:

dateExpr is any expression that evaluates to a date.

MONTH_OF_QUARTER

This function returns the number (between 1 and 3) corresponding to the month in the quarter for a specified date.

Syntax

`MONTH_OF_QUARTER(dateExpr)`

Where:

dateExpr is any expression that evaluates to a date.

MONTHNAME

This function returns the name of the month for a specified date.

Syntax

```
MONTHNAME(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

NOW

This function returns the current timestamp. The `NOW` function is equivalent to the `CURRENT_TIMESTAMP` function.

Syntax

```
NOW()
```

QUARTER_OF_YEAR

This function returns the number (between 1 and 4) corresponding to the quarter of the year for a specified date.

Syntax

```
QUARTER_OF_YEAR(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

SECOND

This function returns the number (between 0 and 59) corresponding to the seconds for a specified time.

Syntax

```
SECOND(timeExpr)
```

Where:

timeExpr is any expression that evaluates to a time.

TIMESTAMPADD

This function adds a specified number of intervals to a specified timestamp, and returns a single timestamp.

In the simplest scenario, this function adds the specified integer value to the appropriate component of the timestamp, based on the interval. Adding a week translates to adding seven days, and adding a quarter translates to adding three months. A negative integer value results in a subtraction (such as going back in time).

An overflow of the specified component (such as more than 60 seconds, 24 hours, 12 months, and so on) necessitates adding an appropriate amount to the next component. For example, when adding to the day component of a timestamp, this function considers overflow and takes into account the number of days in a particular month (including leap years when February has 29 days).

When adding to the month component of a timestamp, this function verifies that the resulting timestamp has enough days for the day component. For example, adding 1 month to 2000-05-31 does not result in 2000-06-31 because June does not have 31 days. This function reduces the day component to the last day of the month, 2000-06-30 in this example.

A similar issue arises when adding to the year component of a timestamp having a month component of February and a day component of 29 (that is, last day of February in a leap year). If the resulting timestamp does not fall on a leap year, the function reduces the day component to 28.

These actions conform to the behavior of Microsoft SQL Server and the native OCI interface for Oracle Database.

Syntax

```
TIMESTAMPADD(interval, intExpr, timestamp)
```

Where:

interval is the specified interval. Valid values are:

- SQL_TSI_SECOND
- SQL_TSI_MINUTE
- SQL_TSI_HOUR
- SQL_TSI_DAY
- SQL_TSI_WEEK
- SQL_TSI_MONTH
- SQL_TSI_QUARTER
- SQL_TSI_YEAR

intExpr is any expression that evaluates to an integer value.

timestamp is any valid timestamp. This value is used as the base in the calculation.

A null integer expression or a null timestamp passed to this function results in a null return value.

Examples

The following query asks for the resulting timestamp when 3 days are added to 2000-02-27 14:30:00. Since February, 2000 is a leap year, the query returns a single timestamp of 2000-03-01 14:30:00.

```
SELECT TIMESTAMPADD(SQL_TSI_DAY, 3, TIMESTAMP'2000-02-27 14:30:00')  
FROM Employee WHERE employeeid = 2;
```

The following query asks for the resulting timestamp when 7 months are added to 1999-07-31 0:0:0. The query returns a single timestamp of 2000-02-29 00:00:00. Notice the reduction of day component to 29 because of the shorter month of February.

```
SELECT TIMESTAMPADD(SQL_TSI_MONTH, 7, TIMESTAMP'1999-07-31 00:00:00')
```

```
FROM Employee WHERE employeeid = 2;
```

The following query asks for the resulting timestamp when 25 minutes are added to 2000-07-31 23:35:00. The query returns a single timestamp of 2000-08-01 00:00:00. Notice the propagation of overflow through the month component.

```
SELECT TIMESTAMPADD(SQL_TSI_MINUTE, 25, 'TIMESTAMP'2000-07-31 23:35:00')
FROM Employee WHERE employeeid = 2;
```

Caution: The `TIMESTAMPADD` function is turned on by default for Microsoft SQL Server, ODBC, IBM DB2, and Oracle databases. Because DB2 and Oracle semantics do not fully support this function, the answers from this function might not match with what the Oracle BI Server computes.

TIMESTAMPDIFF

This function returns the total number of specified intervals between two timestamps.

This function first determines the timestamp component that corresponds to the specified interval parameter, and then looks at the higher order components of both timestamps to calculate the total number of intervals for each timestamp. For example, if the specified interval corresponds to the month component, the function calculates the total number of months for each timestamp by adding the month component and twelve times the year component. Then the function subtracts the first timestamp's total number of intervals from the second timestamp's total number of intervals.

The `TIMESTAMPDIFF` function rounds up to the next integer whenever fractional intervals represent a crossing of an interval boundary. For example, the difference in years between 1999-12-31 and 2000-01-01 is one year because the fractional year represents a crossing from one year to the next (such as 1999 to 2000). By contrast, the difference between 1999-01-01 and 1999-12-31 is zero years because the fractional interval falls entirely within a particular year (that is, 1999). Microsoft SQL Server exhibits the same rounding behavior, but IBM DB2 does not; it always rounds down.

When calculating the difference in weeks, the function calculates the difference in days and divides by seven before rounding. Additionally, the function takes into account how the parameter `FIRST_DAY_OF_THE_WEEK` has been configured in the `NQSConfig.INI` file. For example, with Sunday as the start of the week, the difference in weeks between 2000-07-06 (a Thursday) and 2000-07-10 (the following Monday) results in a value of 1 week. With Tuesday as the start of the week, however, the function would return zero weeks since the fractional interval falls entirely within a particular week. When calculating the difference in quarters, the function calculates the difference in months and divides by three before rounding.

The Oracle BI Server pushes down the `TIMESTAMPADD` and `TIMESTAMPDIFF` functions to Microsoft SQL Server, Oracle Database, IBM DB2, and ODBC databases by default.

Syntax

```
TIMESTAMPDIFF(interval, timestamp1, timestamp2)
```

Where:

interval is the specified interval. Valid values are:

- `SQL_TSI_SECOND`

- SQL_TSI_MINUTE
- SQL_TSI_HOUR
- SQL_TSI_DAY
- SQL_TSI_WEEK
- SQL_TSI_MONTH
- SQL_TSI_QUARTER
- SQL_TSI_YEAR

timestamp1 and *timestamp2* are any valid timestamps.

A null timestamp parameter passed to this function results in a null return value.

Example

The following example query asks for a difference in days between timestamps 1998-07-31 23:35:00 and 2000-04-01 14:24:00. It returns a value of 610. Notice that the leap year in 2000 results in an additional day.

```
SELECT TIMESTAMPDIFF
(SQL_TSI_DAY, TIMESTAMP'1998-07-31 23:35:00',TIMESTAMP'2000-04-01 14:24:00')
FROM Employee WHERE employeeid = 2;
```

Caution: The `TIMESTAMPDIFF` function is turned on by default for Microsoft SQL Server, ODBC, IBM DB2, and Oracle databases. Because DB2 and Oracle semantics do not fully support this function, the answers from this function might not match with what the Oracle BI Server computes.

WEEK_OF_QUARTER

This function returns a number (between 1 and 13) corresponding to the week of the quarter for the specified date.

Syntax

```
WEEK_OF_QUARTER(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

WEEK_OF_YEAR

This function returns a number (between 1 and 53) corresponding to the week of the year for the specified date.

Syntax

```
WEEK_OF_YEAR(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

YEAR

This function returns the year for the specified date.

Syntax

```
YEAR(dateExpr)
```

Where:

dateExpr is any expression that evaluates to a date.

Conversion Functions

The conversion functions convert a value from one form to another. You can also use the VALUEOF function in a filter to reference the value of an Oracle BI system variable. Functions include:

- [CAST](#)
- [CHOOSE](#)
- [IFNULL](#)
- [INDEXCOL](#)
- [TO_DATETIME](#)
- [VALUEOF](#)

CAST

This function changes the data type of an expression or a null literal to another data type. For example, you can cast a customer_name (a data type of Char or Varchar) or birthdate (a datetime literal). The following are the supported data types to which the value can be changed:

```
CHARACTER, VARCHAR, INTEGER, FLOAT, SMALLINT, DOUBLE PRECISION, DATE, TIME,
TIMESTAMP, BIT, BIT VARYING
```

Depending on the source data type, some destination types are not supported. For example, if the source data type is a BIT string, the destination data type must be a character string or another BIT string.

Use CAST to change to a DATE data type. Do not use TO_DATE.

The following describes unique characteristics of the CHAR and VARCHAR data types:

- **Casting to a CHAR data type.** You must use a size parameter. If you do not add a size parameter, a default of 30 is added. Syntax options appear in the following list:

- The recommended syntax is:

```
CAST(expr|NULL AS CHAR(n))
```

For example:

```
CAST(companyname AS CHAR(35))
```

- You can also use the following syntax:

```
CAST(expr|NULL AS data_type)
```

For example:

```
CAST(companyname AS CHAR)
```

Note: If you use this syntax, the Oracle BI Server explicitly converts and stores as `CAST(expr | NULL AS CHAR(30))`

- **Casting to a VARCHAR data type.** You must use a size parameter. If you omit the size parameter, you cannot save the change.

Examples

```
CAST(hiredate AS CHAR(40)) FROM employee
```

```
SELECT CAST(hiredate AS VARCHAR(40)), CAST(age AS double precision), CAST(hiredate AS timestamp), CAST(age AS integer) FROM employee
```

```
CAST("db"."table"."col" AS date)
```

CHOOSE

This function takes an arbitrary number of parameters and returns the first item in the list that the user has permission to see. However, administrators must model the column permissions in the Administration Tool to enable this behavior. See ["INDEXCOL"](#) for an alternate method.

Syntax

```
CHOOSE(expr1, expr2, ..., exprN)
```

For example, a single query can be written to return security-based revenue numbers for the entire organization. The function could look like the following:

```
CHOOSE(L1-Revenue, L2-Revenue, L3-Revenue, L4-Revenue)
```

If the user issuing this function has access to the column L1-Revenue, then that column value would be returned. If the user does not have visibility to the column L1-Revenue but does have visibility to L2-Revenue, then L2-Revenue is returned.

IFNULL

This function tests if an expression evaluates to a null value, and if it does, assigns the specified value to the expression.

Syntax

```
IFNULL(expr, value)
```

Where:

expr is the expression to evaluate.

value is the value to assign if the expression evaluates to a null value.

INDEXCOL

This function can use external information to return the appropriate column for the logged-in user to see. The Oracle BI Server handles this function in the following ways:

- **ODBC Procedures.** `NQSGetLevelDrillability` and `NQSGenerateDrillDownQuery` return the context-specific drill-down information based on the expression translated from `INDEXCOL`. This applies to both `INDEXCOL` expressions specified in the Logical SQL query and `INDEXCOL` expressions specified in a derived logical column.
- **Query Log and cache.** The Logical SQL query with `INDEXCOL` function appears in the SQL string in the query log. But the logical request does not show the `INDEXCOL` function because the Oracle BI Server translates `INDEXCOL` to one of the expressions in its expression list in the logical request generator.
The query cache uses the resulting translated expression for cache hit detection.
- **Usage Tracking.** Usage tracking inserts the Logical SQL query string with the `INDEXCOL` function.
- **Security.** As long as the user has the privileges to access the columns in the expression translated from `INDEXCOL`, then the query executes.
When the first argument to `INDEXCOL` is a session variable and if a default expression is expected to be returned even if the initialization block fails, then you should set a default value for the session variable. Otherwise, the query fails because the session variable has no value definition.

Syntax

```
INDEXCOL(integer_literal, expr_list)
```

Where:

expr_list equals the following:

```
expr1 [, expr_list ]
```

The `INDEXCOL` function takes in an integer literal value as its first argument, followed by a variable length expression list and translates to a single expression from the expression list. The literal value is the 0-based index of the expression in the expression list to translate to. Consider the following expression:

```
INDEXCOL(integer_literal, expr1, expr2, ...)
```

If the literal value is 0, the above expression is equivalent to *expr1*. If the literal value is 1, then the value is equivalent to *expr2*, and so on.

The primary use case for `INDEXCOL` is for the first argument to contain a session variable. Specifying a constant literal would result in `INDEXCOL` always choosing the same expression.

Example With Hierarchy Levels

Company ABC has a geography dimension with the hierarchy Country, State, City. The CEO can access the Country level down to the City level, and the sales manager can access the State and City levels, and the sales people can only access the City level. [Table C-2](#) shows the back-end database for Company ABC.

Table C-2 IndexCol Example of Back-End Database

USER_NAME	TITLE	GEO_LEVEL	CURRENCY	CURRENCY_COL
Bob	CEO	0	US Dollars	0
Harriet	Sales Manager	1	Japanese Yen	1

Table C-2 (Cont.) IndexCol Example of Back-End Database

USER_NAME	TITLE	GEO_LEVEL	CURRENCY	CURRENCY_COL
Jackson	Sales Manager	1	Japanese Yen	1
Mike	Sales Person	2	Japanese Yen	1
Jim	Sales Person	2	US Dollars	0

The following steps illustrate one way to create a single query where each user sees the top level to which they have access:

- The administrator creates a new session variable called `GEOGRAPHY_LEVEL` that is populated by the following initialization block: `SELECT GEO_LEVEL from T where USER_NAME = ':USER'`.

This assumes that the Oracle BI Server instance has the same user names.

- Using `SELECT INDEXCOL(VALUEOF(NQ_SESSION.GEOGRAPHY_LEVEL), Country, State, City), Revenue FROM Sales`, the following occurs:
 - Bob logs in and `INDEXCOL` translates to the Country column because the `GEOGRAPHY_LEVEL` session variable is 0. He gets the same result and can drill down on Country to State as if he had used `SELECT Country, Revenue FROM Sales`.
 - Jackson logs in and `INDEXCOL` translates to the State column because the `GEOGRAPHY_LEVEL` session variable for Jackson is 1. He gets the same result and can drill down on State to City as if he had used `SELECT State, Revenue FROM Sales`.
 - Mike logs in and `INDEXCOL` translates to the City column because the `GEOGRAPHY_LEVEL` session variable for Mike is 2. He gets the same result and cannot drill down on City as if he had used `SELECT City, Revenue FROM Sales`.

TO_DATETIME

This function converts string literals of date`Time` format to a `DateTime` data type.

Syntax

```
TO_DATETIME('string1', 'DateTime_formatting_string')
```

Where:

string1 is the string literal you want to convert

DateTime_formatting_string is the `DateTime` format you want to use, such as `yyyy.mm.dd hh:mi:ss`. For this argument, `yyyy` represents year, `mm` represents month, `dd` represents day, `hh` represents hour, `mi` represents minutes, and `ss` represents seconds.

Examples

```
SELECT TO_DATETIME('2009-03-03 01:01:00', 'yyyy-mm-dd hh:mi:ss') FROM
snowflakesales
```

```
SELECT TO_DATETIME('2009.03.03 01:01:00', 'yyyy.mm.dd hh:mi:ss') FROM
snowflakesales
```

VALUEOF

Use the `VALUEOF` function to reference the value of a repository variable. Repository variables are defined using the Administration Tool. You can use the `VALUEOF` function both in Expression Builder in the Administration Tool, and when you edit the SQL statements for an analysis from the Advanced tab of the Analysis editor in Answers.

Syntax

Variables should be used as arguments of the `VALUEOF` function. Refer to static repository variables by name. Note that variable names are case sensitive. For example, to use the value of a static repository variables named `prime_begin` and `prime_end`:

```
CASE WHEN "Hour" >= VALUEOF("prime_begin")AND "Hour" < VALUEOF("prime_end") THEN
'Prime Time' WHEN ... ELSE...END
```

You must refer to a dynamic repository variable by its fully qualified name. If you are using a dynamic repository variable, the names of the initialization block and the repository variable must be enclosed in double quotes ("), separated by a period, and contained within parentheses. For example, to use the value of a dynamic repository variable named `REGION` contained in an initialization block named `Region Security`, use the following syntax:

```
SalesSubjectArea.Customer.Region = VALUEOF("Region Security"."REGION")
```

The names of session variables must be preceded by `NQ_SESSION`, separated by a period, and contained within parentheses, including the `NQ_SESSION` portion. If the variable name contains a space, enclose the name in double quotes ("). For example, to use the value of a session variable named `REGION`, use the following syntax in Expression Builder or a filter:

```
"SalesSubjectArea"."Customer"."Region" = VALUEOF(NQ_SESSION.REGION)
```

Although using initialization block names with session variables (just as with other repository variables) may work, you should use `NQ_SESSION`. `NQ_SESSION` acts like a wildcard that matches all initialization block names. This lets you change the structure of the initialization blocks in a localized manner without impacting requests.

Database Functions

Users and administrators can create requests by directly calling database functions from either Oracle BI Answers, or by using a logical column (in the logical table source) within the metadata repository. Key uses for these functions include the ability to pass through expressions to get advanced calculations, as well as the ability to access custom written functions or procedures on the underlying database.

Support for database functions does not currently extend across all multidimensional sources. Also, you cannot use these functions with XML data sources.

Functions include:

- [EVALUATE](#)
- [EVALUATE_ANALYTIC](#)
- [EVALUATE_AGGR](#)
- [EVALUATE_PREDICATE](#)

EVALUATE

This function passes the specified database function with optional referenced columns as parameters to the back-end data source for evaluation. This function is intended for scalar calculations, and is useful when you want to use a specialized database function that is not supported by the Oracle BI Server, but that is understood by the underlying data source.

The embedded database function may require one or more columns. These columns are referenced by %1 ... %N within the function. The actual columns must be listed after the function.

Syntax

```
EVALUATE('db_function(%1...%N)' [AS data_type] [, column1, columnN])
```

Where:

db_function is any valid database function understood by the underlying data source.

data_type is an optional parameter that specifies the data type of the return result. Use this parameter whenever the return data type cannot be reliably predicted from the input arguments. However, do not use this parameter for type casting; if the function needs to return a particular data type, add an explicit cast. You can typically omit this parameter when the database-specific function has a return type not supported by the Oracle BI Server, but is used to generate an intermediate result that does not need to be returned to the Oracle BI Server.

column1 through *columnN* is an optional, comma-delimited list of columns.

Examples

This example shows an embedded database function.

```
SELECT EVALUATE('instr(%1, %2)', address, 'Foster City') FROM employees
```

Examples Using EVALUATE_AGGREGATE and EVALUATE to Leverage Unique Essbase Functions

The following examples use the EVALUATE_AGGREGATE and EVALUATE functions. Note that expressions are applied to columns in the logical table source that refers to the physical cube.

Use EVALUATE_AGGREGATE to implement custom aggregations. For example, you may want to compare overall regional profit to profits for the top three products in the region. You can define a new measure to represent the profits for top three products resulting in the Logical SQL statement:

```
SELECT Region, Profit, EVALUATE_AGGREGATE('SUM(TopCount(%1.members, 3, %2), %3)', Products, Profit, Profit) Top_3_prod_Profit FROM SampleBasic
```

The Oracle BI Server generates the following expression for the custom aggregation:

```
member [Measures].[MS1] AS
'SUM(Topcount([Product].Generations(6).members,3,[Measures].[Profit]),[Measures].[Profit])'
```

Use the EVALUATE function on projected dimensions to implement scalar functions that are computed post-aggregation. EVALUATE may change the grain of the query, if its definition makes explicit references to dimensions (or attributes) that are not in the query.

For example, if you would like to see the Profits for the top five products ranked by Sales sold in a Region, after creating the applicable measure, the resulting Logical SQL statement is as follows

```
SELECT Region, EVALUATE('TopCount(%1.members, 5, %2)' as VARCHAR(20), Products,
Sales), Profits
FROM SampleBasic
```

The Oracle BI Server generates the following expression to retrieve the top five products:

```
set [Evaluate0] as
'{Topcount([Product].Generations(6).members,5,[Measures].[Sales])}'
```

EVALUATE_ANALYTIC

This function passes the specified database analytic function with optional referenced columns as parameters to the back-end data source for evaluation.

The embedded database function may require one or more columns. These columns are referenced by %1 ... %N within the function. The actual columns must be listed after the function.

Syntax

```
EVALUATE_ANALYTIC('db_function(%1...%N)' [AS data_type] [, column1, columnN])
```

Where:

db_function is any valid database analytic function understood by the underlying data source.

data_type is an optional parameter that specifies the data type of the return result. Use this parameter whenever the return data type cannot be reliably predicted from the input arguments. However, do not use this parameter for type casting; if the function needs to return a particular data type, add an explicit cast. You can typically omit this parameter when the database-specific analytic function has a return type not supported by the Oracle BI Server, but is used to generate an intermediate result that does not need to be returned to the Oracle BI Server.

column1 through *columnN* is an optional, comma-delimited list of columns.

Examples

This example shows an embedded database analytic function.

```
EVALUATE_ANALYTIC('dense_rank() over(order by %1 )' AS INT,sales.revenue)
```

If the preceding example needs to return a double, then an explicit cast should be added, as follows:

```
CAST(EVALUATE_ANALYTIC('Rank(%1.dimension.currentmember, %2.members)',
"Foodmart93"."Time"."Month" as Double)
```

EVALUATE_AGGR

This function passes the specified database function with optional referenced columns as parameters to the back-end data source for evaluation. This function is intended for aggregate functions with a GROUP BY clause.

The embedded database function may require one or more columns. These columns are referenced by %1 ... %N within the function. The actual columns must be listed after the function.

Syntax

```
EVALUATE_AGGR('db_agg_function(%1...%N)' [AS data_type] [, column1, columnN])
```

Where:

db_agg_function is any valid aggregate database function understood by the underlying data source.

data_type is an optional parameter that specifies the data type of the return result. Use this parameter whenever the return data type cannot be reliably predicted from the input arguments. However, do not use this parameter for type casting; if the function needs to return a particular data type, add an explicit cast. You can typically omit this parameter when the database-specific function has a return type not supported by the Oracle BI Server, but is used to generate an intermediate result that does not need to be returned to the Oracle BI Server.

column1 through *columnN* is an optional, comma-delimited list of columns.

Example

```
EVALUATE_AGGR('REGR_SLOPE(%1, %2)', sales.quantity, market.marketkey)
```

EVALUATE_PREDICATE

This function passes the specified database function with optional referenced columns as parameters to the back-end data source for evaluation. This function is intended for functions with a return type of Boolean.

The embedded database function may require one or more columns. These columns are referenced by %1 ... %N within the function. The actual columns must be listed after the function.

Note that EVALUATE_PREDICATE is not supported for use with Essbase data sources.

Syntax

```
EVALUATE_PREDICATE('db_function(%1...%N)', [, column1, columnN])
```

Where:

db_function is any valid database function with a return type of Boolean that is understood by the underlying data source.

column1 through *columnN* is an optional, comma-delimited list of columns.

If you want to model a database function for comparison purposes, you should not use EVALUATE_PREDICATE. Instead, use EVALUATE and put the comparison outside the function. For example, do *not* use EVALUATE_PREDICATE as follows:

```
EVALUATE_PREDICATE('dense_rank() over (order by 1% ) < 5', sales.revenue)
```

Instead, use EVALUATE, as follows:

```
EVALUATE('dense_rank() over (order by 1% ) ', sales.revenue) < 5
```

Example

```
SELECT year, Sales AS DOUBLE, CAST(EVALUATE('OLAP_EXPRESSION(%1, 'LAG(units_cube_sales, 1, time, time LEVELREL time_levelrel)')', OLAP_CALC) AS DOUBLE) FROM
```



```
"Global".Time, "Global"."Facts - sales" WHERE EVALUATE_PREDICATE('OLAP_
CONDITION(%1, 'LIMIT time KEEP ''1'', ''2'', ''3'', ''4'' ')
=1', OLAP_CALC) ORDER BY year;
```

Hierarchy Navigation Functions

The hierarchy navigation functions enable you to identify relationships between members of hierarchies. The hierarchy navigation functions include:

- [ISANCESTOR](#)
- [ISCHILD](#)
- [ISDESCENDANT](#)
- [ISLEAF](#)
- [ISPARENT](#)
- [ISROOT](#)

The [ISLEAF](#) function applies to both level-based and parent-child hierarchies, while the other functions apply only to parent-child hierarchies.

See [Chapter 9, "Working with Logical Dimensions"](#) for information about level-based and parent-child hierarchies, including information about creating parent-child relationship tables (closure tables) for relational sources.

ISANCESTOR

The [ISANCESTOR](#) function enables you to find the ancestors of a member of a parent-child hierarchy, either all the ancestors of a member, or the ancestors at a specified hierarchical distance from the member.

Each member of the parent-child hierarchy is compared with the specified member to determine if it is an ancestor. The [ISANCESTOR](#) function returns the Boolean value True for each ancestor of the specified member, else it returns False.

You can use the [ISANCESTOR](#) function in a query both within CASE statements and in WHERE clause conditions.

You can use the [ISANCESTOR](#) function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

```
ISANCESTOR(pc_presentation_hierarchy, member_identifier [, distance])
```

Where:

pc_presentation_hierarchy identifies the fully qualified parent-child presentation hierarchy, as follows:

```
"subject_area"."presentation_table"."pc_presentation_hierarchy"
```

The qualification term ("subject_area".) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

member_identifier is the string or numeric literal that identifies the member in *pc_presentation_hierarchy*. The type of literal depends on the data type of the dimension level keys.

distance (optional) is a positive integer that identifies the distance from the specified member to the parent-child hierarchy level at which to search for ancestors.

By default, if *distance* is not specified, the ISANCESTOR function searches the current parent-child level containing *member_identifier* and all levels above.

Business Model and Mapping Layer Syntax

```
ISANCESTOR(logical_dimension, member_identifier [, distance])
```

Where:

logical_dimension identifies the fully qualified dimension containing the parent-child hierarchy, as follows:

```
"business_model"."dimension_name"
```

The qualification term ("business_model" .) is optional unless there are multiple dimensions with the same name in different business models.

member_identifier is the string or numeric literal that identifies the member in *logical_dimension*. The type of literal depends on the data type of the dimension level keys.

distance (optional) is a positive integer that identifies the distance from the specified member to the parent-child hierarchy level at which to search for ancestors.

By default, if *distance* is not specified, the ISANCESTOR function searches the current parent-child level containing *member_identifier* and all levels above.

Example

The following example selects all the ancestor employees of the employee Joe in a parent-child hierarchy. The returned list includes the employee Joe.

```
SELECT emp_name
FROM "employees"
WHERE ISANCESTOR("employees"."emp_hierarchy", 'Joe')
```

ISCHILD

The ISCHILD function enables you to find the children of a member of a parent-child hierarchy, that is, all the members that are one hierarchical level below the specified member.

Note: The ISCHILD function is the same as the ISDESCENDANT function with a distance parameter of 1.

The ISCHILD function returns the Boolean value True for each child of the specified member, else it returns False.

You can use the ISCHILD function in a query both within CASE statements and in WHERE clause conditions.

You can use the ISCHILD function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

```
ISCHILD(pc_presentation_hierarchy, member_identifier)
```

Where:

pc_presentation_hierarchy identifies the fully qualified parent-child presentation hierarchy, as follows:

```
"subject_area"."presentation_table"."pc_presentation_hierarchy"
```

The qualification term ("subject_area" .) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

member_identifier is the string or numeric literal that identifies the member in *pc_presentation_hierarchy*. The type of literal depends on the data type of the dimension level-keys.

Business Model and Mapping Layer Syntax

```
ISCHILD(logical_dimension, member_identifier)
```

Where:

logical_dimension identifies the fully qualified dimension containing the parent-child hierarchy, as follows:

```
"business_model"."dimension_name"
```

The qualification term ("business_model" .) is optional unless there are multiple dimensions with the same name in different business models.

member_identifier is the string or numeric literal that identifies the member in *logical_dimension*. The type of literal depends on the data type of the dimension level keys.

Example

The following example selects all the children of the employee Joe in a parent-child hierarchy.

```
SELECT emp_name
FROM "employees"
WHERE ISCHILD("employees"."emp_hierarchy", 'Joe')
```

ISDESCENDANT

The `ISDESCENDANT` function enables you to find the descendants of a member of a parent-child hierarchy, either all the descendants of a member, or the descendants at a specified hierarchical distance from the member.

Each member of the parent-child hierarchy is compared with the specified member to determine if it is a descendant. The `ISDESCENDANT` function returns the Boolean value `True` for each descendant of the specified member, else it returns `False`.

You can use the `ISDESCENDANT` function in a query both within `CASE` statements and in `WHERE` clause conditions.

You can use the `ISDESCENDANT` function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

```
ISDESCENDANT(pc_presentation_hierarchy, member_identifier [, distance])
```

Where:

pc_presentation_hierarchy identifies the fully qualified parent-child presentation hierarchy, as follows:

```
"subject_area"."presentation_table"."pc_presentation_hierarchy"
```

The qualification term ("subject_area" .) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

member_identifier is the string or numeric literal that identifies the member in *pc_presentation_hierarchy*. The type of literal depends on the data type of the dimension level-keys.

distance (*optional*) is a positive integer, that identifies the distance from the specified member to the parent-child hierarchy level at which to search for descendants.

By default, if *distance* is not specified, the ISDESCENDANT function searches the current parent-child level containing *member_identifier* and all levels below.

Business Model and Mapping Layer Syntax

```
ISDESCENDANT(logical_dimension, member_identifier [, distance])
```

Where:

logical_dimension identifies the fully qualified dimension containing the parent-child hierarchy, as follows:

```
"business_model"."dimension_name"
```

The qualification term ("business_model" .) is optional unless there are multiple dimensions with the same name in different business models.

member_identifier is the string or numeric literal that identifies the member in *logical_dimension*. The type of literal depends on the data type of the dimension level-keys.

distance (*optional*) is a positive integer that identifies the distance from the specified member to the parent-child hierarchy level at which to search for descendants.

By default, if *distance* is not specified, the ISDESCENDANT function searches the current parent-child level containing *member_identifier* and all levels below.

Example

The following example selects all the descendant employees of the employee Joe in a parent-child hierarchy. The returned list includes the employee Joe.

```
SELECT emp_name
FROM "employees"
WHERE ISDESCENDANT("employees"."emp_hierarchy", 'Joe')
```

ISLEAF

The ISLEAF function applies to both level-based and parent-child hierarchies. For both types of hierarchy, a leaf member is defined as a member that has no child members.

Each member of the hierarchy is examined to determine if it is a leaf member. The ISLEAF function returns the Boolean value True for each leaf member, else it returns False.

You can use the ISLEAF function in a query both within CASE statements and in WHERE clause conditions.

You can use the ISLEAF function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

ISLEAF(*presentation_hierarchy*)

Where:

presentation_hierarchy identifies the fully qualified presentation hierarchy, either level-based or parent-child, as follows:

"subject_area"."presentation_table"."presentation_hierarchy"

The qualification term ("subject_area".) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

Business Model and Mapping Layer Syntax

ISLEAF(*logical_dimension*)

Where:

logical_dimension identifies the fully qualified dimension containing the hierarchy you want to navigate, either level-based or parent-child, as follows:

"business_model"."dimension_name"

The qualification term ("business_model".) is optional unless there are multiple dimensions with the same name in different business models.

Example

The following example selects all the employees in a hierarchy that are leaf members, that is, the employees who have no members below them in the hierarchy.

```
SELECT emp_name
FROM "employees"
WHERE ISLEAF("employees"."emp_hierarchy")
```

ISPARENT

The ISPARENT function enables you to find the parents of a member of a parent-child hierarchy, that is, all the members that are one hierarchical level above the specified member.

Note: The ISPARENT function is the same as the ISANCESTOR function with a distance parameter of 1.

The ISPARENT function returns the Boolean value True for each parent of the specified member, else it returns False.

You can use the `ISPARENT` function in a query both within `CASE` statements and in `WHERE` clause conditions.

You can use the `ISPARENT` function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

```
ISPARENT(pc_presentation_hierarchy, member_identifier)
```

Where:

pc_presentation_hierarchy identifies the fully qualified parent-child presentation hierarchy, as follows:

```
"subject_area"."presentation_table"."pc_presentation_hierarchy"
```

The qualification term ("`subject_area`".) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

member_identifier is the string or numeric literal that identifies the member in *pc_presentation_hierarchy*. The type of literal depends on the data type of the dimension level keys.

Business Model and Mapping Layer Syntax

```
ISPARENT(logical_dimension, member_identifier)
```

Where:

logical_dimension identifies the fully qualified dimension containing the parent-child hierarchy, as follows:

```
"business_model"."dimension_name"
```

The qualification term ("`business_model`".) is optional unless there are multiple dimensions with the same name in different business models.

member_identifier is the string or numeric literal that identifies the member in *logical_dimension*. The type of literal depends on the data type of the dimension level keys.

Example

The following example selects all the parents of the employee Joe in a parent-child hierarchy.

```
SELECT emp_name
FROM "employees"
WHERE ISPARENT("employees"."emp_hierarchy", 'Joe')
```

ISROOT

A presentation hierarchy member is defined as a root member if it has no ancestors above it in a parent-child presentation hierarchy.

Each member of the parent-child hierarchy is examined to determine if it is a root member. The `ISROOT` function returns the Boolean value `True` for each root member, else it returns `False`.

You can use the ISROOT function in a query both within CASE statements and in WHERE clause conditions.

You can use the ISROOT function in both Presentation layer queries, and in the Business Model and Mapping layer (for example, when creating a derived column). Note that the syntax of the function depends on where you are using it.

Presentation Layer Syntax

```
ISROOT(pc_presentation_hierarchy)
```

Where:

pc_presentation_hierarchy identifies the fully qualified parent-child presentation hierarchy, as follows:

```
"subject_area"."presentation_table"."pc_presentation_hierarchy"
```

The qualification term ("subject_area" .) is optional unless there are multiple presentation tables or presentation hierarchies with the same name in different subject areas.

Business Model and Mapping Layer Syntax

```
ISROOT(logical_dimension)
```

Where:

logical_dimension identifies the fully qualified dimension containing the parent-child hierarchy, as follows:

```
"business_model"."dimension_name"
```

The qualification term ("business_model" .) is optional unless there are multiple dimensions with the same name in different business models.

Example

The following example selects all the employees in a parent-child hierarchy that are root members, that is, the employees who have no ancestors above them in the hierarchy.

```
SELECT emp_name
FROM "employees"
WHERE ISROOT("employees"."emp_hierarchy")
```

System Functions

The system functions return values relating to the session. Functions include:

- [USER](#)
- [DATABASE](#)

USER

This function returns the user name for the Oracle BI repository to which you are logged on.

Syntax

```
USER()
```

DATABASE

This function returns the name of the default subject area.

Syntax

```
DATABASE ( )
```

Merge Rules

When you use the Merge Repository Wizard in the Administration Tool, sophisticated rules determine how objects are merged. Some decisions about merging objects are made automatically by the wizard, while other decisions appear as prompts in the Define Merge Strategy screen. This appendix describes some of the merge rules and behavior of the Merge Repository Wizard.

There are three types of merges:

- Full merges are typically used during development processes, when there are two different repositories that need to be merged. The Administration Tool provides a three-way merge feature that lets you merge two repositories that have both been derived from a third, original repository. Full merges can also be used to import objects from one repository into another. See ["Performing Full Repository Merges"](#) for more information.
- Patch merges are used when you are applying the differential between two versions of the same repository. For example, you might want to use a patch merge to apply changes from the development version of a repository to your production repository, or to upgrade your Oracle BI Applications repository. See ["Performing Patch Merges"](#) for more information.
- Multiuser development merges are used when you are checking in projects using a multiuser development environment. See ["About the Multiuser Development Merge Process"](#) for more information.

This appendix contains the following topics:

- [General Merge Rules and Behavior](#)
- [Special Merge Algorithms for Logical Table Sources and Other Objects](#)

General Merge Rules and Behavior

The merge process typically involves three versions of an Oracle BI repository: the **original** repository, **modified** repository, and **current** repository. The original repository is the original unedited file (the parent repository), while the modified and current repository are the two changed files you want to merge. The current repository is the one currently open in the Administration Tool.

The original, modified, and current repository may mean different things, depending on your situation. For example:

- In a development-to-production scenario, you have an original parent file, a current file that contains the latest development changes, and a modified file that is the deployed copy of the original file.

- In an Oracle BI Applications repository upgrade scenario, the current file is the latest version of the repository shipped by Oracle, and the original file is the original repository shipped by Oracle. The modified file is the file that contains the customizations you made to the original file.

Note that patch merge can be used with both of these situations. In a patch merge, you open the current file and select the original file, then generate the patch. To apply the patch, you open the modified file and select the original file, then apply the patch. See ["Performing Patch Merges"](#) for more information.

Regardless of which merge scenario you want to perform, and regardless of the merge type (full, patch, or multiuser development merge), the following general rules are applied:

- It is assumed that you generally want to keep the changes in the modified repository. For example, if an object is added to or deleted from the modified repository, the object is added or deleted without prompting.
- If an object is added to or deleted from the current repository, the Merge Repository Wizard asks whether you want to keep the changes.

In general, the Merge Repository Wizard tries to ensure that you have the minimum set of objects necessary to service your queries. During a merge, there might be objects introduced by the current repository that are not needed in the merged repository. To address this issue, the Merge Repository Wizard asks whether new Presentation layer objects in the current repository are needed in the final merged result. If you choose to keep the new presentation objects, all the dependent logical and physical objects are added as well. However, if you choose not to keep the new presentation objects, then the dependent logical and physical objects are not kept, because no queries will require the use of these objects. The Merge Repository Wizard discards these objects to ensure that the merged repository does not get populated with unused objects.

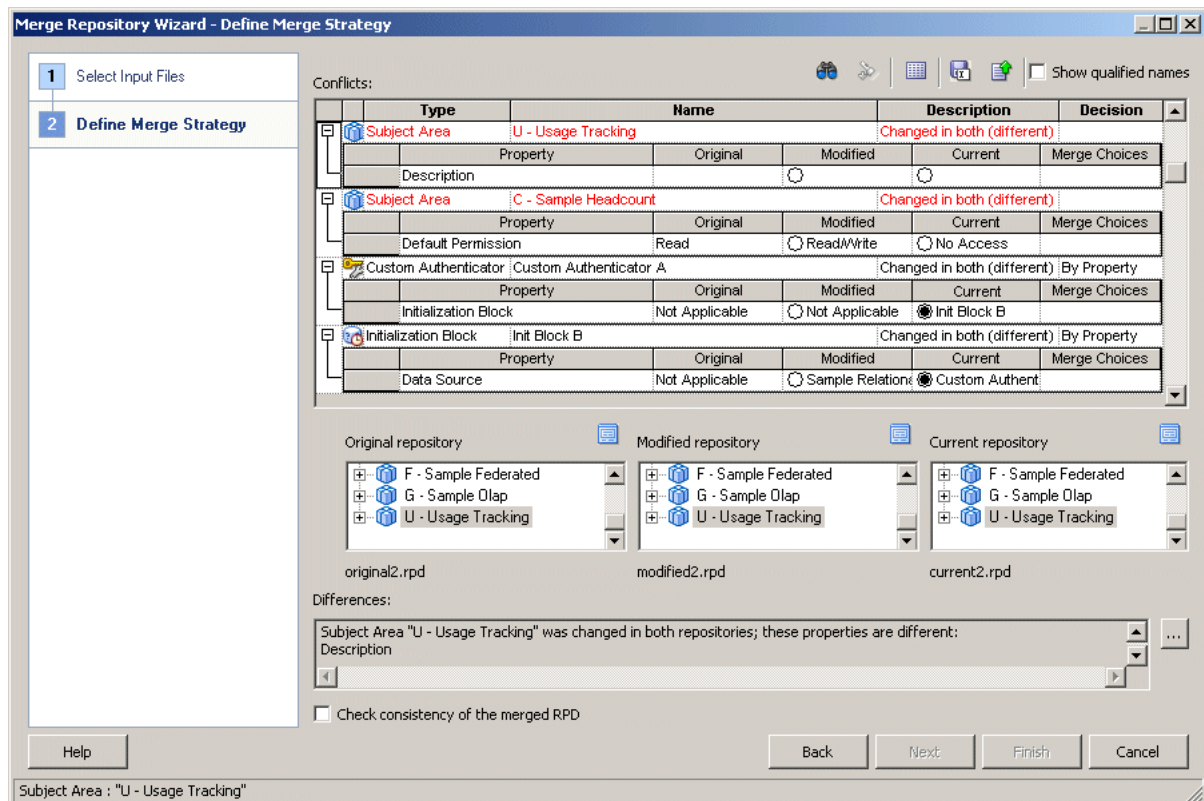
- If an object is added to or deleted from both repositories, the object is added or deleted without prompting. If the same object was added with slight differences in its properties, the Merge Repository Wizard asks which version of the properties you want to keep.
- If an object has been modified only in the current repository, or only in the modified repository, the change is kept. If the same object is modified in both the current and modified repository, and the changes are different, then there is a merge conflict. When conflicts occur, the Merge Repository Wizard asks you to choose which change you want to keep.
- Making a decision about one object can determine a whole set of decisions, depending on the object relationships involved. For example, if you choose to keep a presentation column that has been added to the current repository, then the associated presentation table and subject area must also be kept, along with the logical column, physical column, and other associated objects upon which it is based. Alternatively, if you choose not to keep a subject area that has been added to the current repository, then you are not prompted to choose whether to keep its associated objects. Adding a join may require the addition of its base tables, while changing an expression may cause physical columns to be added.
- Object relationships can be interconnected through their properties. In addition to strings and numbers, the internal value of a property can be other repository objects. Because of this, a change to one object might cause a corresponding change to an interrelated object.

For example, assume you change the data source of Init Block B from a connection pool to Custom Authenticator A. In addition to the data source property change to the initialization block object, a corresponding property change occurs in the custom authenticator object (because the value of the initialization block property for Custom Authenticator A is now Init Block B).

Because the decisions made for these properties must be synchronized, if you select **Current** as the decision for the data source property of Init Block B, then the decision for the initialization block property of Custom Authenticator A will also be **Current**. See Figure D-1 shows what this example looks like in the Merge Repository Wizard.

In the Merge Repository Wizard, any decisions that require user input are displayed on the Define Merge Strategy screen. Figure D-1 shows the Define Merge Strategy screen.

Figure D-1 Define Merge Strategy Screen



Some other rules are dependent on the type of merge you want to perform. For example, if you are performing a patch merge to upgrade a repository, you want to retain the security settings and database feature table changes in your customized (modified) repository. If you are performing a multiuser development (MUD) merge, however, you do not want to retain security settings and database feature table changes, to prevent developers from overwriting passwords and other important objects in the master repository. Because of this, changes to security settings and database features are not retained when you perform a MUD merge.

To change security settings or database features in a multiuser development environment, you must edit the master repository directly. To do this, remove the master repository from the multiuser development directory, edit it in offline mode, then move it back.

Special Merge Algorithms for Logical Table Sources and Other Objects

In addition to the general rules governing how objects are merged and which situations require prompting, there are special rules for certain types of objects and situations.

This section contains the following topics:

- [Merging Objects that Use the Vector Merge Algorithm](#)
- [Merging Logical Table Sources](#)
- [Merging Security Filters](#)
- [Inferring the Use Logical Column Property for Presentation Columns](#)
- [Merging Aliases](#)

Merging Objects that Use the Vector Merge Algorithm

Some objects, such as levels, application roles, and object permissions, use a vector merge algorithm that determines the parent/child relationships between objects.

Objects that use the vector merge algorithm include:

- Levels in a dimension, levels associated with a logical column, and child levels
- Dimensions and tables in a logical display folder
- Aggregation content in a logical table source
- Security objects like user and application role membership and permissions
- Initialization block LDAP server settings and execution precedence

The Oracle BI Server determines the initial state of object relationships in each repository during the merge process. For example, the following list shows the different possibilities for object permissions and how they relate to users and application roles:

- M - Missing. The application role, user, or object is not present in the repository.
- D - Default. Permissions are inherited from the parent application role.
- Y - Yes. The permission is explicitly granted to the user or application role.
- N - No. The permission is explicitly denied to the user or application role.

The Merge Repository Wizard determines the appropriate relationship for the merged repository depending on the state of the object permission relationships in each repository. For example:

- For an original repository with a result of Y, a modified repository with a result of N, and a current repository with a result of M, the Merge Repository Wizard determines a result of N for the merged repository.
- For an original repository with a result of N, a modified repository with a result of Y, and a current repository with a result of M, the Merge Repository Wizard determines a result of Y for the merged repository.

[Example D-1](#) provides a detailed explanation of how object relationships are merged for application role objects.

Example D-1 Vector Merge Example: Merging Application Roles

The following list shows the different possibilities for user and application role relationships:

- M - Missing. The application role or user is not present in the repository.
- Y - Yes. The application role or user is a member of the application role.
- N - No. The application role or user is not a member of the application role.

Table D-1 shows the merged result for different combinations of object relationships in the merging repositories.

Table D-1 Results for Merging Application Roles Based on Object Relationships

Original Repository	Modified Repository	Current Repository	Result
M	M	M	N ¹
M	M	Y	Y
M	M	N	N
M	Y	M	Y
M	Y	Y	Impossible ²
M	Y	N	Impossible
M	N	M	N
M	N	Y	Impossible
M	N	N	Impossible
Y	M	M	Y
Y	M	Y	Y
Y	M	N	N
Y	Y	M	Y
Y	Y	Y	Y
Y	Y	N	N
Y	N	M	N
Y	N	Y	N
Y	N	N	N
N	M	M	N
N	M	Y	Y
N	M	N	N
N	Y	M	Y
N	Y	Y	Y
N	Y	N	Y
N	N	M	N
N	N	Y	Y
N	N	N	N

¹ This situation can happen if neither the application role nor the user are present in the original repository, but the user is present in the modified repository and the application role is present in the current repository. In this case, no membership can be assumed.

² M in original implies that either the user or application role is not present. The missing object added in both cannot be considered the same object.

Merging Logical Table Sources

Special rules govern how to merge column mappings in logical table source objects. Each column mapping is merged individually. For each column, if the mapping has changed in either the modified or current repository, the change is kept. If the mapping has changed in both repositories, the Oracle BI Server attempts to merge the mappings automatically.

Note that the deletion of a column is not considered to be a change in its mapping. If a column is not present in the modified repository, then the mapping in the current repository is used instead.

If there are differences in aggregation content, then the aggregation content specified by level has priority. In other words, if the aggregation content in one repository is by level and the aggregation content in another repository is by column, then the aggregation content by level is retained.

Merging Security Filters

If a filter for an application role has changed in only one repository, then the change is kept. If the filter has changed in both repositories, the Oracle BI Server attempts to merge the filters automatically.

If an object is required for merging a particular filter (such as a presentation column) and is not present, then that filter is considered invalid and does not appear in the merged repository. Note, however, that this rule does not apply to variables. If a variable is required for merging a particular filter, the Oracle BI Server ensures that the variable is retained in the merged repository.

Inferring the Use Logical Column Property for Presentation Columns

Presentation columns have both a Name property and a Use Logical Column Name property. In some cases, these properties can come into conflict. For example, [Table D-2](#) shows a scenario where this situation could occur.

Table D-2 *Conflicting Presentation Column Name and Use Logical Column Name Properties*

Repository	Presentation Column Name	Logical Column Name	Use Logical Column Name
Original	Sales	GroupSales	No
Current	Sales	Sales	Yes
Modified	GroupSales	GroupSales	Yes

If the regular merge rules for the objects in [Table D-2](#) are applied, the merged repository contains a presentation column called GroupSales and a logical column called Sales, with the Use Logical Column Name property set to Yes. However, this result is incorrect, because the name of the presentation column is different from the name of the logical column.

To avoid this situation, the Oracle BI Server infers the value of the Use Logical Column Name property. Using this logic, the merged repository for the example in [Table D-2](#) has a presentation column called GroupSales, a logical column called Sales, and a Use Logical Column Name property set to No.

Merging Aliases

During the full merge process, users are not prompted to make decisions about aliases. Aliases from the current and modified repositories are merged automatically.

In multiuser development merges, however, users are prompted to choose whether to keep aliases from the current repository, keep aliases from the modified repository, or merge choices to keep aliases from both repositories.

Also note the following:

- If object names change because of the merge process, then the previous names are added as aliases.
- Any aliases that are not associated with presentation objects are deleted.

Deleting Unwanted Objects from the Repository

If your repository contains many objects that you do not need, you can use the command-line pruning utility to delete the unwanted objects. This appendix explains how to the pruning utility works and how to use it.

This appendix contains the following topics:

- [About the Object Pruning Utility](#)
- [Using the Object Pruning Utility](#)
- [Deletion Rules for the Object Pruning Utility](#)

About the Object Pruning Utility

If you have a large number of extraneous or unwanted objects in your repository, you can delete the unwanted objects using the `prunerpd` command-line utility. You can use `prunerpd` on both Windows and UNIX systems.

You can delete unwanted repository objects such as databases, tables, columns, initialization blocks, and variables. However, note that the pruning utility does not remove objects from the Oracle BI Presentation Catalog.

Deleting objects from the repository has a cascading effect. For example, if a physical column is deleted, then any mapped logical columns are deleted, as well as any associated presentation columns. See "[Deletion Rules for the Object Pruning Utility](#)" for more information.

Using the Object Pruning Utility

You must first create the input file that contains the list of repository objects to be deleted. Then, you must run the utility at the command line, passing the input file as an argument

This section contains the following topics:

- [Creating the Input File](#)
- [Running the Prune Utility](#)

Creating the Input File

The prune utility accepts the list of repository objects you want to delete as a text file. The utility can accept multiple input files at a time. The syntax rules for the input file are shown in [Table E-1](#).

Note: Object names in the input file must match the fully qualified name that is used in the repository. Wildcards (such as "*" and "?") are not supported in the object name.

Table E-1 Syntax Rules for Input File

Object Type	Example	Action
Database	D "Paint"	Deletes the database named "Paint."
Table	<ul style="list-style-type: none"> ▪ T "W_AGREE_D" ▪ T "DB"."Catalog"."Schema"."Table" 	<ul style="list-style-type: none"> ▪ Deletes the table or alias named "W_AGREE_D" from the Physical layer. ▪ Deletes the table or alias named "Table" from the schema named "Schema," contained in the catalog named "Catalog," located in the database named "DB," from the Physical layer.
Column	C "W_AGREE_MD"."AGREE_CD"	Deletes the column named "AGREE_CD" located in a table or alias named "W_AGREE_D" from the Physical layer.
Initialization block	I "External Metadata Strings"	Deletes the initialization block named "External Metadata Strings."
Variable	V CURR_USER	Deletes the variable named "CURR_USER."

For example, a text file that contains instructions to delete a database named "Stock Quotes" and a physical column named "S_NQ_ACCT"."USER_NAME" would include the following entry:

```
D "Stock Quotes" C "S_NQ_ACCT"."USER_NAME"
```

Use white space as a delimiter in the input file (a single space, tab, or multiple spaces).

Running the Prune Utility

Before running `prunerpd`, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX systems) to launch a command prompt that is initialized to your Oracle instance. You can find this utility in:

```
ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup
```

Then, run `prunerpd` from the resulting command prompt window with the desired options. The prune utility accepts the following parameters:

-s

Provide the location of the source repository file.

-p

Provide the repository password.

The password argument is optional. If you do not provide a password argument, you are prompted to enter a password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide a password argument either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.

-f

Provide the input file name (in text format) that contains the list of repository objects to be removed. Separate multiple file names by spaces. Enclose spaces within a filename with double quotes (" ").

-o

Provide the name and location of the output repository file, also known as the pruned repository.

-l

Provide the name and location of the output log file. All actions performed on the repository are written to this file, including descriptions.

-e

Provide the name and location of the error log file. The pruning utility writes exceptions and errors to this log.

-8

Input is in UTF-8 format.

Example E-1 Prune Utility Usage Example

```
prunerpd -s C:/OBI/Server/Repository/BIApps.rpd
-f "C:/Remove Oracle EBS Objects.txt"
-o "C:/OBI/Server/Repository/BIApps Pruned.rpd"
-l "C:/temp/BIApps Pruning.log" -e "C:/temp/BIApps Pruning.err"
Give password: my_repos_password
```

Log File

The log file includes a list of actions performed on the repository file and is in XML format. Other messages, such as progress indicators, are sent to the standard output stream.

Error Log File

The error log file includes a list of errors encountered while processing the contents of the input file and is in XML format. Other errors are sent to the standard output error stream.

Deletion Rules for the Object Pruning Utility

Deleting repository objects has a cascading effect. This section describes the deletion rules.

Physical Layer Rules

- If a physical column or a table is deleted, then all of the affected keys, foreign keys, and complex joins are deleted as well. The internal obsolete attribute definition (attr defn) that links a logical column to a physical column is also removed.
- Empty schemas, catalogs, and databases are removed.
- If a table is deleted, then all its columns are deleted.

Logical Table Rules

- If a regular column (not an aggregate or derived column) is not mapped in any logical source, then it is deleted. The keys, including the level key and the logical key, are also removed.

- If the source column for a derived column or its referenced variable is deleted (corrupted), then the column is removed.
- If an aggregate rule or override aggregate rule for an aggregate column is corrupted (due to a logical column deletion), then the column is removed.
- If a logical table is removed (because its underlying physical table was deleted), then the keys, foreign keys, logical joins, sources, and source folder are removed.
- If a logical table source does not have any valid mapping, then it is deleted.
- If a logical table source is retained, but its aggregate content or filters are corrupted, then the corresponding expressions are set to null. The join specification is also removed.
- If a logical table, dimension, or business model is empty (contains no meaningful child), then it is deleted.

Presentation Layer Rules

- If a logical column is removed (because its underlying physical column was deleted), then any corresponding presentation columns are removed.
- If a presentation table or subject area does not contain children, then it is removed.

Security Rules

- If a security filter for a user or application role becomes corrupt due to deletion, then the filter is removed. If all filters are removed for a user or application role, then the internal privilege object is deleted.
- Even if all filters for an application role are deleted, the application role is still maintained.
- To remove an application role from the repository, you must explicitly delete it. See *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition* for information about deleting application roles.

Variable Rules

- Initialization blocks are deleted if the underlying connection pool is deleted.
- Repository and session variables are deleted if the associated initialization blocks are deleted.
- If a session variable is deleted and its parent initialization block does not contain variables, then the initialization block is removed.
- If an initialization block is deleted, then its variables are removed.

Marketing Rules

- Qualified list items are deleted if the associated cache catalog, GUID column, or qualified column is deleted.
- Target levels are deleted if the associated catalog (Segmentation Catalog name) is deleted.
- List catalogs are deleted if the associated catalog, table, or column is deleted.
- Conforming dimensions are deleted if the associated catalog, table, or column is deleted.

Exchanging Metadata with Databases to Enhance Query Performance

If your organization has installed either Oracle Database or IBM DB2, then you can use these databases to enhance the data warehouse performance and functionality of queries that run on the Oracle BI Server.

This appendix contains the following topics:

- [About Exchanging Metadata with Databases](#)
- [Generating the Import File](#)
- [Using Materialized Views in the Oracle Database with Oracle Business Intelligence](#)
- [Using IBM DB2 Cube Views with Oracle Business Intelligence](#)

About Exchanging Metadata with Databases

By exchanging Oracle Business Intelligence metadata from the Oracle BI Server with your Oracle Database or IBM DB2 database, you enable the database to accelerate the performance of data warehouse queries.

You use the Oracle BI Server utility `sametaexport` to exchange the metadata. When you run this utility to generate cube views for DB2, the utility is called the DB2 Cube Views Generator. When you run this utility to generate metadata for Oracle Database, the utility is called the Oracle Database Metadata Generator.

The Oracle BI Server export utility works with the following tools:

- In the Oracle Database, the SQL Access Advisor creates materialized views and index recommendations on optimizing performance. Note that in database releases prior to 10g, this feature is called the Oracle Database Summary Advisor.
- In the IBM DB2 database, IBM DB2 Cube Views creates materialized query tables (MQTs).

The `sametaexport` utility generates the information necessary for the SQL Access Advisor or IBM DB2 Cube Views tool to preaggregate the relational data and improve query performance.

Generating the Import File

Both the Oracle Database Metadata Generator and the DB2 Cube Views Generator create the files that are needed to import metadata from the Oracle BI Server into the SQL Access Advisor or an IBM DB2 database.

This section contains the following topics that are common to the two generators:

- [Running the Generator](#)
- [About the Metadata Input File](#)
- [About the Output Files](#)
- [Troubleshooting Errors from the Generator](#)
- [Metadata Conversion Rules and Error Messages](#)

Running the Generator

The Oracle Database Metadata Generator and the DB2 Cube Views Generator are invoked from the command line or embedded in a batch file. The command-line executable is named `sametaexport`.

Before running the utility, you must first run `bi-init.cmd` (or `bi-init.sh` on UNIX) to launch a command prompt or shell window that is initialized to your Oracle instance. You can find this utility in:

`ORACLE_INSTANCE/bifoundation/OracleBIApplication/coreapplication/setup`

Then, run `sametaexport` with the desired options, as follows:

```
sametaexport -r "PathAndRepositoryFileName" [-p repository_password]
-f "InputFileNameAndPath" [options]
```

[Table 18–3](#) contains descriptions of the parameters in the command-line executable file.

Table 18–3 Parameters for `sametaexport`

Parameter	Definition	Additional Information
-r	Repository file name and full path	Quotation marks are required for the file name and path only if the file path is in long format or has spaces. Use the full path if the file is not in the current directory.
-p	Repository password	The password for the given repository. The password argument is optional. If you do not provide a password argument, you are prompted to enter a password when you run the command. To minimize the risk of security breaches, Oracle recommends that you do not provide a password argument either on the command line or in scripts. Note that the password argument is supported for backward compatibility only, and will be removed in a future release.
-f	Input file name and full path	Quotation marks are required for the file name and path only if the file path is in long format or has spaces. Use the full path if the file is not in the current directory. You specify input files so that you do not have to type all the required information at the command line, and so that you can type international characters. See " About the Metadata Input File " for more information.

You can include some additional parameters in the input file or at the command line to change various defaults for the Oracle Database Metadata Generator and the DB2 Cube Views Generator. Parameters specified in the input file take precedence over parameters specified at the command line. You must include these parameters only if you want to change the default values.

Table 18–4 and Table 18–5 describe these optional parameters.

Table 18–4 Optional Parameters and Defaults for the Oracle Database Metadata Generator

Parameter Definition	Additional Information	Input File Usage Example	Command Line Usage Example
Use schema name from RPD	When set to YES, the table schema names are used as they are used in the repository. The default value is YES.	USE_SCHEMA_NAME_FROM_RPD = NO	-schemafrom rpd NO
Default schema name	The default schema name is used as the table schema name if the value of <code>-schemafromrpd</code> is set to NO, or if the repository schema name cannot be determined. The default value is STEEBEL.	DEFAULT_SCHEMA_NAME = ORACLE	-defaultschema ORACLE
Oracle schema name	The metadata from Oracle Database Metadata Generator is created under this schema. The default value is STEEBEL.	ORA_DIM_SCHEMA_NAME = ORACLE	-orclschema ORACLE
Logging enabled	Indicates whether to keep a log of the metadata export process. Valid values are ON, OFF, and DEBUG. The default value is ON.	LOGGING = DEBUG	-logging DEBUG
Log file name	The path to the log file. If you provide an invalid path, an error occurs. If you do not provide this parameter, the default log file path is used. The default path is: <code>ORACLE_INSTANCE\diagnostics\logs\OracleBIServerComponent\coreapplication_obisn\OraDimExp.log</code>	LOG_FILE_NAME = C:\bea_default\instances\instance1\diagnostics\logs\generator\logfile.log	-logfile C:\bea_default\instances\instance1\diagnostics\logs\generator\logfile.log

Table 18–5 Optional Parameters and Defaults for the DB2 Cube Views Generator

Parameter Definition	Additional Information	Input File Usage Example	Command Line Usage Example
Distinct count supported	When set to YES, allows measure containing the DISTINCT_COUNT aggregation to be exported. The recommended setting and default value is NO.	DISTINCT_COUNT_SUPPORTED = YES	-distinct YES
Statistical functions supported	When set to YES, allows measures containing the aggregation STDDEV to be exported. The recommended setting and default value is NO.	STATISTICAL_FUNCTIONS_SUPPORTED = YES	-stat YES
Use schema name	When set to YES, the Cube Views metadata attributes have columns from tables under a schema name, which are then specified in the parameters. When set to NO, the schema names for these tables are empty. The default value is YES.	USE_SCHEMA_NAME = NO	-useschema NO

Table 18–5 (Cont.) Optional Parameters and Defaults for the DB2 Cube Views Generator

Parameter Definition	Additional Information	Input File Usage Example	Command Line Usage Example
Use schema name from RPD	When set to YES, the table schema names are used as they are used in the repository. The default value is YES.	USE_SCHEMA_NAME_FROM_RPD = NO	-schemafromrpd NO
Default schema name	The default schema name is used as the table schema name if the value of <code>-schemafromrpd</code> is set to NO, or if the repository schema name cannot be determined. The default value is <code>SIEBEL</code> .	DEFAULT_SCHEMA_NAME = ORACLE	-defaultschema ORACLE
Cube views schema name	The name of the schema under which the Cube Views metadata is created. The default value is <code>SIEBEL</code> .	CUBE_VIEWS_SCHEMA_NAME = ORACLE	-cubeschema ORACLE
Log file name	The path to the log file. If you provide an invalid path, an error occurs. If you do not provide this parameter, the default log file path is used. The default path is: <code>ORACLE_INSTANCE\diagnostics\logs\OracleBIServerComponent\coreapplication_obisn\CubeViews.log</code>	LOG_FILE_NAME = C:\bea_default\instances\instance1\diagnostics\logs\generator\logfile.log	-logfile C:\bea_default\instances\instance1\diagnostics\logs\generator\logfile.log
Log failures	When set to YES, the log file lists the metadata that was invalidated under a certain rule. The default value is YES.	LOG_FAILURES = NO	-logfail NO
Log success	When set to YES, the log file lists the metadata that has been checked under each rule and has passed the check. The default value is NO.	LOG_SUCCESS = YES	-logsuccess YES

About the Metadata Input File

The input file is a text file that contains the parameters that are described in [Table 18–6](#).

Table 18–6 Cube Metadata Input File Parameters

Input File Name	Description
BUSINESS_MODEL	The name of the business model in the logical layer of the Oracle Business Intelligence repository that contains the metadata to export. If the business model is not found in the repository, then an error message is displayed. You can only specify one business model name in the input file. To generate metadata for multiple business models, create another input file and run the Oracle Database Metadata Generator or DB2 Cube Views Generator again.

Table 18–6 (Cont.) Cube Metadata Input File Parameters

Input File Name	Description
PHYSICAL_DATABASE	The name of the database in the physical layer of the Oracle Business Intelligence repository that contains the metadata to export. When the business model derives from multiple databases, then it eliminates metadata from all databases other than the one specified here. When the physical database is not found in the repository, an error message is displayed.
RUN_AS_USER	The user name of the database user whose visibility must be duplicated for the metadata export. This parameter cannot be empty. This user must exist as a user reference in the repository.
OUTPUT_FOLDER	The full path and file name of the folder to which the SQL file will be written. If the folder does not exist when you run the Oracle Database Metadata Generator, then it will be created. See "About the Output Files" for more information.

The following text shows a sample metadata input file:

```
BUSINESS_MODEL = "1 - Sample App"
PHYSICAL_DATABASE = "1 - Sample App Data"
RUN_AS_USER = "Administrator"
OUTPUT_FOLDER = "C:\OracleBI"
```

About the Output Files

Each Generator creates different types of output files, as described in the following list:

- **Oracle Database Metadata Generator:** Generates a SQL file that is encoded in UTF-8 and stored in the specified output folder. The file name is based on the name of the business model you specified in the input file, such as my_business_model.sql.
- **DB2 Cube Views Generator:** Generates the following files in the specified output folder:
 - XML file (encoded in UTF8). One XML file is created for the specified business model. It contains all objects that were converted to cubes. Additionally, objects in the repository will be mapped to similar objects in the IBM Cube Views metadata. See ["Conversion Rules for IBM DB2 Databases"](#) for a list of objects that will not be converted.
 - The name of the XML file matches the business model name, without spaces, followed by the XML extension (for example, SalesResults.xml).
 - A SQL file that contains the alias generation DLL. A SQL file is created for the specified business model only if aliases exist in the physical layer databases that are referenced in the business model. The alias file contains SQL commands that will create the aliases in the DB2 database. The name of the SQL file matches the business model name, without spaces, followed by the SQL extension (for example, SalesResults-alias.sql).

Troubleshooting Errors from the Generator

Error messages indicate that the Generator was unable to complete some or all of its tasks. After starting the Generator, you might observe the following error messages:

- Unable to write to Log file: *log_file_name*.

The log file specified in the input file or at the command line might contain the wrong path, the user might not have write permissions to that folder, or the disk could be out-of-space.

- `Run_as_user, user_name, is invalid.`
The user name is incorrect.
- `Repository, repository_name.rpd, is invalid or corrupt.`
The repository name might be incorrect, it might not exist in the given path, or the user might not have permission to read it.
- `Physical Database, database_name, is invalid.`
The physical database name does not match a valid physical database object in the repository.
- `Business Model, model_name, is invalid.`
The business model name does not match a valid business model object in the repository.
- `Authentication information provided is invalid.`
The repository password provided at the command line is incorrect.
- `Path: "path_name" is invalid.`
The path or file name is incorrect, or the current user does not have read access.

Metadata Conversion Rules and Error Messages

When the Generator creates the output files, it also maps the metadata objects in the Oracle Business Intelligence repository to similar objects in the metadata of the Oracle Database or the IBM DB2 database.

This section explains the rules used to identify Oracle Business Intelligence metadata that cannot be translated (converted) into either SQL or XML format. These rules are necessary because Oracle Database and IBM Cube Views do not support some of the metadata constructs that are allowed by Oracle Business Intelligence.

Dimensional metadata in the SQL or XML file will be generated at the logical fact table source level. If a logical fact table source has an invalid logical dimension table source, then the logical dimension table source will be invalidated. If the logical fact table source is invalid, then all the logical dimension table sources that are mapped to it will also be invalidated. Invalid Oracle Business Intelligence repository metadata elements will not be converted to cubes in the SQL or XML file.

When a rule is violated, the Generator writes the error messages and the metadata that violated the rule to the log file.

Conversion Rules for Oracle Databases

The following list provides the rules for converting Oracle Business Intelligence metadata into objects in the Oracle Database:

- Attributes that contain expressions in the logical table cannot be exported.
- Tables joined using complex joins are not considered.
- Tables that are opaque views are not considered.
- Columns used as part of a key in one level cannot be used as part of another level key.

Oracle Database prohibits the use of columns as keys in multiple levels. This prohibition requires the Oracle Database Metadata Generator to eliminate one of the two joins, usually the join that is encountered first. Therefore, the other joins are lost, which prevents them from being exported.

Conversion Rules for IBM DB2 Databases

Table 18–7 lists the rules used to validate Oracle Business Intelligence repository metadata elements, error messages that are written to the log file if the rule is violated, and an explanation of what caused the rule violation. The error messages help you determine why a particular Oracle Business Intelligence metadata object was not exported to the XML file.

Table 18–7 Validation Rules for Metadata Elements

Rule	Message	Explanation
ComplexJoin FactsRule	[Fact Logical Table Source]Complex Physical Joins not supported %qn has a complex Join %qn between Physical Tables %qn and %qn	If the physical fact tables are connected through complex joins, then the join is not supported. A complex join is defined as any join between two tables that do not have a foreign key relationship.
ComplexJoin DimsRule	[Dimension Logical Table Source]Complex Physical Joins not supported %qn has a complex Join %qn between Physical Tables %qn and %qn	If the dimension physical tables are connected through a complex join, then that join is not supported.
ComplexJoin FactDimRule	[Fact Logical Table Source -> Dimension Logical Table Source] Complex Physical Joins not supported. %qn has a complex Join %qn between Physical Tables %qn and %qn.	If a dimension physical table and a fact physical table are connected through a complex join, then that join is not supported and the dimension table source is invalidated.
OpaqueView FactRule	[Fact Logical table Source] Physical SQL Select Statements not supported. %qn uses the SQL Select Statement %qn.	When the physical fact table is generated by a SQL select statement, the logical fact table source that contains the table is invalidated. All logical dimension table sources connected to this logical fact table source are also invalidated. This construct allows subquery processing.
OpaqueView DimRule	[Dimension Logical table Source] Physical SQL Select Statements not supported. %qn uses the SQL Select Statement %qn.	When a physical dimension table is generated by a SQL select statement, the logical dimension table source containing that table is invalidated.
OuterJoinFactRule	[Fact Logical Table Source] Physical Outer Joins not supported. %qn has an outer join %qn between physical tables %qn and %qn.	If the logical fact table source has an outer join mapping, then that logical fact table source is invalidated and all logical dimension table sources mapped to this source will also be invalidated.
OuterJoinDimRule	[Dimension Logical Table Source] Physical Outer Joins not supported. %qn has an outer join %qn between physical tables %qn and %qn.	If the logical dimension table source has an outer join mapping, then that logical dimension table source is invalidated.

Table 18–7 (Cont.) Validation Rules for Metadata Elements

Rule	Message	Explanation
WhereClauseFactRule	[Fact Logical Table Source] WHERE clauses are not supported. %qn has a where condition %s.	If the fact table source uses a WHERE clause to filter the data that is loaded, then this table source is invalidated.
WhereClauseDimRule	[Dimension Logical Table Source] WHERE clauses are not supported. %qn has a where condition %s.	If the dimension table source uses a WHERE clause to filter the data that is loaded, then this table source is invalidated.
TwoJoinFactDimRule	[Fact Logical Table Source -> Dimension Logical Table Source] Multiple Joins between sources not supported. %qn and %qn have at least the following joins : %qn, %qn.	If a physical fact table is mapped to two dimension tables from the same dimension source (if the fact table is not exclusively mapped to the most detailed table in the table source), then the dimension table source is invalidated.
HiddenManyRule	[Fact Logical Table Source -> Dimension Logical Table Source] Join between (physical or logical?) fact and dimension is not on the most detailed table. %qn between %qn and %qn is not on the most detailed table %qn (Join name, facttable, dimtable).	This is related to the TwoJoinFactDimRule. If the fact table is joined to a dimension table that is not the most detailed table in the table source, then the dimension table source is invalidated.
ComplexMeasureRule	[Column] Complex Aggregation Rules not supported. %qn uses an aggregation rule of %s which is not supported.	The supported aggregations are typically SUM, COUNT, AVG, MIN, MAX, STDDEV, COUNTDISTINCT, and COUNT.
CountDistinctMeasureRule	[Column] COUNT-DISTINCT Aggregation Rule not supported. %qn uses an aggregation rule of %s which is not supported.	COUNTDISTINCT aggregation is not supported for this particular column.
InvalidColumnLevelRule	[Level] Some columns that are part of the Primary Level Key are invalid. %qn has %qn as part of its primary key, when %qn has already been marked invalid.	The level key for this level has one or more columns that are invalid.
VariableBasedColumnRule	[Logical Table Source -> Column] Column uses a Variable in the Expression Column %qn uses a variable in its mapping.	The logical column uses repository and session variables in the expression.
OneFactToManyDimRule	[Fact Logical Table Source -> Dimension Logical Table Source] There must be a unique join path between the most detailed tables in the (logical or physical?) fact and the dimension. No join paths found between %qn and %qn (both physical table names). Found at least the following join paths: (%qn->%qn...), (%qn->%qn...)	Same as in TwoJoinFactDimRule or HiddenManyRule.

Table 18–7 (Cont.) Validation Rules for Metadata Elements

Rule	Message	Explanation
ManyMDTinFactRule	[Fact Logical Table Source] Fact Logical Table Source must have a unique most detailed table. %qn has at least the following most detailed tables : %qn,%qn.	A fact that has more than one table that is the most detailed table.
NoMeasureFactRule	[Fact Logical Table Source] Fact Logical Table Source does not have any Measures. %qn does not have any deployable measures.	A fact table does not have any measures because all the measures have been invalidated.
NoInactiveFactRule	[Fact Logical Table Source] Fact Logical Table Source is not marked Active.	A fact source is not active.
NoInactiveDimensionRule	[Dimension Logical Table Source] Dimension Logical Table Source is not marked Active.	A dimension source is not active.
NoAttributeInFactRule	[Fact Logical Table Source -> Column] Attribute found in Fact. %qn in a fact source %qn does not have an aggregation rule.	No attributes in the fact source.
NoMeasureInDimensionRule	[Dimension Logical Table Source -> Column] Measure found in Dimension. %qn in a dimension source %qn has an aggregation rule.	No measures in the dimension source.
VisibleColumnsAttrRule	[Column] -> The run_as_user does not have visibility to this Logical Column. %qn is not accessible to the run_as_user %qn due to visibility rules.	A column does not have visibility for this user.
VisibleColumnsMeasureRule	[Column] -> The run_as_user does not have visibility to this Logical Column. %qn is not accessible to the run_as_user %qn due to visibility rules.	A column does not have visibility for this user.
MultiplePrimaryKeysDimensionRule	[Dimension Logical Table Source] A Join uses an alternate key in the Dimension Logical Table Source. %qn between %qn and %qn in %qn uses the alternate key %qn.	A dimension physical table can contain only one primary key. It is joined to another dimension physical table using a different unique key and that join is invalid. IBM Cube Views does not accept any unique keys to be used for foreign joins and always requires the primary key.
MultiplePrimaryKeysFactRule	[Dimension Logical Table Source] A Join uses an alternate key in the Dimension Logical Table Source. %qn between %qn and %qn in %qn uses the alternate key %qn.	A fact physical table can contain only one primary key. It is joined to another fact physical table using a different unique key and that join is invalid. IBM Cube Views does not accept any unique keys to be used for foreign joins and always requires the primary key.

Table 18–7 (Cont.) Validation Rules for Metadata Elements

Rule	Message	Explanation
MultiplePrimaryKeysFactDimRule	[Fact Logical Table Source -> Dim Logical Table Source] A Join uses an alternate key between the Logical Table sources. %qn between %qn and %qn for sources %qn and %qn uses the alternate key %qn.	A fact physical table can contain only one primary key. It is joined to a dimension physical table using a different unique key and is invalid. IBM Cube Views does not accept any unique keys to be used for foreign joins and always requires the primary key.
NotDB2ExpressionAttributeRule	[Dimension Logical Table Source -> Column] The Column contains an Expression not supported. %qn has expression %s which is not supported.	The attribute contains an expression not supported by IBM Cube Views. This includes metadata expressions that use DateTime functions (for example, CURRENT_DATE).
NotDB2ExpressionMeasureRule	[Fact Logical Table Source -> Column] The Column contains an Expression not supported. %qn has expression %s which is not supported.	A measure contains an expression not supported by IBM Cube Views. This includes metadata expressions that use DateTime functions (for example, CURRENT_DATE).
NoAttributeDimRule	[Dimension Logical Table Source] Dimension Logical Table Source does not have any attributes visible to the run_as_user. %qn can not be queried by user %qn since none of its attributes are visible.	A dimension does not have any attributes.

Using Materialized Views in the Oracle Database with Oracle Business Intelligence

This section explains how to export metadata from Oracle Business Intelligence into the SQL Access Advisor and create materialized views using the Oracle Database Metadata Generator.

This section contains the following topics:

- [About Using the SQL Access Advisor with Materialized Views](#)
- [Deploying Metadata for Oracle Database](#)

About Using the SQL Access Advisor with Materialized Views

This feature enhances the data warehouse performance and functionality of a database. It enables the SQL Access Advisor to store metadata about the logical relationships of the data that resides in the database. Additionally, it accelerates data warehouse queries by using more efficient Oracle materialized views. These materialized views preaggregate the relational data and improve query performance. Once the metadata is stored in the SQL Access Advisor, the database administrator can optimize the database objects and improve query performance.

When processing queries, Oracle Database routes queries to tables that hold materialized views when possible. Because these tables of materialized views are smaller than the underlying base tables and the data has been pre aggregated, the queries that are rerouted to them might run faster.

Oracle Database Metadata Generator works as a metadata bridge to convert the Oracle Business Intelligence proprietary metadata into a SQL file that contains PL/SQL

commands to generate dimensions in the SQL Access Advisor. After converting metadata into a SQL file, you use a tool such as SQL*Plus to import the translated metadata into the SQL Access Advisor and store it in metadata catalog tables. After importing the metadata, you create materialized views, which are used by to optimize incoming application queries.

You can use this feature with Oracle Database 9i and higher. See "[System Requirements and Certification](#)" for information about platform compatibility.

Note that in database releases prior to 10g, the SQL Access Advisor was called the Oracle Database Summary Advisor and was documented in *Oracle9i Data Warehousing Guide*.

Deploying Metadata for Oracle Database

Become familiar with the Oracle Database and its tools before attempting to deploy metadata in the Oracle Database. For more information, see "SQL Access Advisor" in *Oracle Database Performance Tuning Guide*.

Ensure that you complete the steps in "[Running the Generator](#)" before deploying metadata. To deploy cube metadata, perform the tasks described in the following sections:

- [Executing the SQL File for Oracle Database](#)
- [Defining Constraints for the Existence of Joins](#)
- [Creating the Query Workload](#)
- [Creating Materialized Views](#)

Executing the SQL File for Oracle Database

Before executing the SQL file for importing into the SQL Access Advisor, ensure that you are familiar with Oracle Database import tools. See the Oracle Database documentation set for information.

Use a tool such as SQL*Plus to execute the SQL file that the Oracle Database Metadata Generator generated. You might see error messages if the dimensions already exist or if the database schema differs from that in the RPD file. When the script executes successfully, you can see the dimensions that were created by using the database web console or the Oracle Enterprise Manager Database Control. In the Oracle Enterprise Manager Database Control, expand the following nodes: Network, Databases, database-name, Warehouse, Summary Management, Dimensions, System.

After you execute the SQL file, be aware of the following:

- No incremental metadata changes are allowed. Schema changes require that you manually delete cube model metadata in the Oracle Database and convert the Oracle Business Intelligence metadata again. For example, if you must make a change to a dimension in a cube in the Oracle BI repository, you must delete the cube model in the Oracle Database, regenerate the SQL file from the Oracle BI repository, and import it into the SQL Access Advisor.
- You cannot delete metadata using the Oracle Database Metadata Generator. Instead, you must manually delete the cube model using the Oracle Enterprise Manager Database Control.

Defining Constraints for the Existence of Joins

For more information on this topic, see the Oracle Database documentation set.

You must ensure that Oracle Database knows about the joins between the dimension tables and the fact tables. To do so, you create constraints in SQL*Plus or the Oracle Enterprise Manager Database Control. In the Oracle Enterprise Manager Database Control, you select the table on which you must create a constraint, then select the Constraint tab.

You create a different type of constraint for each kind of table, as follows:

- For dimension tables, create a `UNIQUE` key constraint.
- For fact tables, create a `FOREIGN` key constraint and specify the referenced schema and referenced table. In the Constraint Definition area, include the foreign key columns in the fact table and the corresponding unique keys in the dimension table. An attempt to create a foreign key on a fact table can fail if the foreign key column data does not match the unique key column data on the dimension table.

Creating the Query Workload

See the Oracle Database documentation set for detailed information about creating the query workload.

A query workload is a sample set of physical queries to optimize. Before you create the workload, you generate a Trace file with information on the slowest-running queries.

To generate the Trace file:

You can generate the Trace file of the slowest-running queries using a tool that is appropriate to your database version, as described in the following list:

- **Usage Tracking:** Use this capability in Oracle Business Intelligence to log queries and how long they take to run. Long-running Oracle Business Intelligence queries can then be executed as a script and used with the Trace feature in the Oracle Database to capture the Oracle Database SQL code for these queries.
- **Oracle Database Trace:** Use this tool to identify the slowest physical query. You can enable the Trace feature either within Oracle Enterprise Manager Database Control or by entering SQL commands with the `DBMS_MONITOR` package. Once you enable the Trace feature, you use a script to create a Trace file to capture the SQL code for queries in a query workload table.
- **Oracle Enterprise Manager:** Use this tool to track slow-running queries.

Note: The capabilities that are described in the following sections are available in Oracle Database, rather than as part of Oracle Business Intelligence.

To analyze the information in the Trace file:

1. Use the following guidelines when reviewing the Trace file:
 - When you have traced many statements at once, such as in batch processes, quickly discard any statements that have acceptable query execution times. Focus on those statements that take the longest times to execute.
 - Check the Query column for block visits for read consistency, including all query and subquery processing. Inefficient statements are often associated with a large number of block visits. The Current column indicates visits not related to read consistency, including segment headers and blocks that will be updated.

- Check the Disk column for the number of blocks that were read from disk. Because disk reads are slower than memory reads, the value will likely be significantly lower than the sum of the Query and Current columns. If it is not, check for issues with the buffer cache.
 - Locking problems and inefficient PL/SQL loops can lead to high CPU time values even when the number of block visits is low.
 - Watch for multiple parse calls for a single statement, because this indicates a library cache issue.
2. After identifying the problem statements in the file, check the execution plan to learn why each problem statement occurred.

To load queries into the workload:

- After you use the Trace utility to learn the names of the slowest physical queries, insert them into the USER_WORKLOAD table.

Table 18–8 describes the columns of the USER_WORKLOAD table.

- Use INSERT statements to populate the QUERY column with the SQL statements for the slowest physical queries and the OWNER column with the appropriate owner names.

Table 18–8 Columns in USER_WORKLOAD Table

Column	Data Type	Required	Description
QUERY	Any LONG or VARCHAR type (all character types)	YES	SQL statement for the query.
OWNER	VARCHAR2 (30)	YES	User who last executed the query.
APPLICATION	VARCHAR2 (30)	NO	Application name for the query.
FREQUENCY	NUMBER	NO	Number of times that the query was executed.
LASTUSE	DATE	NO	Last date on which the query was executed.
PRIORITY	NUMBER	NO	User-supplied ranking of the query.
RESPONSETIME	NUMBER	NO	Execution time of the query in seconds.
RESULTSIZE	NUMBER	NO	Total number of bytes that the query selected.
SQL_ADDR	NUMBER	NO	Cache address of the query.
SQL_HASH	NUMBER	NO	Cache hash value of the query.

Creating Materialized Views

After you populate the query workload table, use the appropriate tool for the Oracle Database version to create materialized views. In Oracle Database 10g, use the SQL Access Advisor in the Oracle Enterprise Manager Database Control and specify the query workload table that you created.

The SQL Access Advisor generates recommendations on improving the performance of the fact tables that you specify. The SQL Access Advisor displays the SQL code with

which it will create the appropriate materialized views. Before indicating that the SQL Access Advisor should create the materialized views, review the following tips:

- The creation of a materialized view can fail if the SQL code includes a `CAST` statement.
- Ensure that the `CREATE MATERIALIZED VIEW` statement does not specify the same query that you provided as a workload table. If the statement does specify the same query, then the materialized views will likely not reflect the true performance gain. However, if the query is executed frequently, then creating a materialized view might still be worthwhile.
- Add a forward slash (/) to the end of the `CREATE MATERIALIZED VIEW` statement after the SQL statement. Otherwise, the SQL*Plus worksheet will not recognize it as a valid statement.

Note: The SQL Access Advisor can also help determine appropriate indexing schemes.

Using IBM DB2 Cube Views with Oracle Business Intelligence

This section explains how to export metadata from Oracle Business Intelligence into IBM DB2 using the DB2 Cube Views Generator.

This section contains the following topics:

- [About Using IBM DB2 Cube Views with Oracle Business Intelligence](#)
- [Deploying Cube Metadata](#)

About Using IBM DB2 Cube Views with Oracle Business Intelligence

The term IBM DB2 Cube Views is a registered trademark of IBM. See "[System Requirements and Certification](#)" for information about platform compatibility.

This feature enhances the data warehouse performance and functionality of a database. It enables the DB2 database to store metadata about the logical relationships of the data residing in the database. Additionally, it accelerates data warehouse queries by using more efficient DB2 materialized query tables (MQTs). These MQTs preaggregate the relational data and improve query performance.

When processing queries, the DB2 Query Rewrite functionality routes queries to the MQTs when possible. Because these tables are smaller than the underlying base tables and the data has been pre aggregated, the queries that are rerouted to them might run faster.

DB2 Cube Views Generator works as a metadata bridge to convert the Oracle Business Intelligence proprietary metadata into an IBM Cube Views XML file. After converting metadata into an XML file, you use IBM Cube Views to import the translated metadata into the DB2 database and store it in IBM Cube Views metadata catalog tables. After importing the metadata, you use the IBM Optimization Advisor to generate scripts to create materialized query tables (MQT) and their indexes. The deployed MQTs are used by the DB2 Query Reroute Engine to optimize incoming application queries.

DB2 provides an API (implemented as a stored procedure) that passes XML documents as arguments to create, modify, delete, or read the metadata objects. For more information about IBM Cube Views, see the IBM DB2 documentation.

Deploying Cube Metadata

The alias-SQL file generated by the DB2 Cube Views Generator should be executed before importing the XML file. The XML file generated by the DB2 Cube Views Generator contains the cube metadata in XML format. After importing the XML file into your DB2 database, you must create materialized query tables.

Note: It is strongly recommended that you become familiar with IBM Cube Views and its tools before attempting to import the XML file. For more information, see the IBM documentation.

Ensure that you complete the steps in "[Running the Generator](#)" before deploying metadata. To deploy cube metadata, perform the tasks described in the following sections:

- [Executing the Alias-SQL File for IBM Cube Views](#)
- [Importing the XML File](#)
- [Guidelines for Creating Materialized Query Tables \(MQTs\)](#)

Executing the Alias-SQL File for IBM Cube Views

You must execute the alias-SQL file before you import the XML file into your DB2 database. For more information, see the IBM documentation.

The alias-SQL file that is generated by the DB2 Cube Views Generator must be executed by a SQL client on the database where the data warehouse is located. When executed, it creates aliases (synonyms) for tables in the database.

Importing the XML File

After you execute the alias-SQL file, you can import the XML file into the database. For more information, see the IBM documentation.

Note: It is strongly recommended that you become familiar with IBM Cube Views and its tools before attempting to import the XML file. For more information, see the IBM documentation.

You can import this file using the following IBM tools:

- **IBM OLAP Center (recommended).** For more information, see "[Guidelines for Importing the XML File Using the IBM OLAP Center](#)" and the IBM documentation.
- **IBM command-line client utility (db2mdapiclient.exe).** IBM ships this utility with DB2. For more information about using the command-line client utility, see the IBM documentation.
- **IBM DB2 Stored Procedure.** IBM Cube Views provides a SQL-based and XML-based application programming interface (API) that you can use to run a single stored procedure to create, modify, and retrieve metadata objects. For more information, see the IBM documentation.

Guidelines for Importing the XML File Using the IBM OLAP Center Using the IBM OLAP Center, you can import cube metadata into DB2. The IBM OLAP Center provides wizards to help you import the file. For more information, see the IBM documentation.

To import the XML file, use the following guidelines:

- Using the IBM OLAP Center tool, connect to the DB2 database.
- In the Import Wizard, choose the XML file that you want to import.
- If metadata exists that refers to database constructs that are not in the database, then an error message is displayed.
- When the wizard asks for an import option, choose to replace existing objects.
- When you are returned to the IBM OLAP Center, a diagram of the cube model is shown.

Guidelines for Changing Cube Metadata After Importing the XML File After you import the XML file, you might need to perform the following actions:

- Because Oracle OLAP does not store foreign keys as metadata, they will not exist in the converted metadata in the DB2 database. You must use the IBM Referential Integrity Utility for IBM Cube Views to generate foreign key informational constraints. You can obtain this utility on the IBM Web site.
- You might encounter other issues such as foreign key join columns being nullable. You can use the following methods to solve this problem:
 - If data in these columns are not null, then you should convert these columns to not-null columns.
 - If data in these columns are null or you prefer not to convert the column data type even if the column data is not null, then you should modify the cube model using the following guidelines:
 - * In a fact-to-dimension join, you must manually eliminate this dimension object from the converted cube model and create a degenerated dimension object consisting of the foreign key of this join.
 - * In a dimension-to-dimension join, you must manually eliminate the dimension object that represents the primary-key side of the join from the converted cube model and create a degenerated dimension object consisting of the foreign key of this join.
 - * In a fact-to-fact join, you must manually eliminate the fact object that represents the primary-key side of the join from the converted cube model and create a degenerated dimension object consisting of the foreign key of this join.
- No incremental metadata changes will be allowed by the Cube Generator. Schema changes require that you manually delete cube model metadata in the DB2 database and convert the Oracle Business Intelligence metadata again. For example, if you must make a change to a dimension in a cube in the Oracle Business Intelligence metadata repository, then you must delete the cube model in the DB2 database, regenerate the XML file from the Oracle Business Intelligence repository, and import it into the DB2 database.
- You cannot delete metadata using the DB2 Cube Views Generator. Instead, you must manually delete the cube model using the IBM OLAP Center.
- The IBM Statistics tool and IBM Optimization Advisor must be run periodically.

For more information, see the IBM documentation.

Guidelines for Creating Materialized Query Tables (MQTs)

For more information, see the IBM documentation.

After you import the cube metadata into the database, you must run the IBM Optimization Advisor to generate SQL scripts and then execute those scripts to create the MQTs. You must provide certain parameters to the IBM Optimization Advisor to get optimal results from the implementation. The IBM Optimization Advisor wizard analyzes your metadata and recommends how to build summary tables that store and index aggregated data for SQL queries. Running the IBM Optimization Advisor can help you keep the MQTs current. Additionally, you must refresh your database after each ETL.

To create MQTs, use the following guidelines:

- In the IBM OLAP Center, choose the cube model that you want to optimize and open the IBM Optimization Advisor wizard.
- Follow the instructions in the wizard, using the following table as a guide.

When asked for:	Choose:
Summary Tables	Choose Deferred (or Immediate) and provide a tablespace for the tables
Limitations	Choose an appropriate value for the optimization parameters. You should turn on the Data-sampling option.
SQL Scripts	Creation of the scripts needed to run to create the Summary tables. Choose the filename and locations

- When the IBM Optimization Advisor closes, you must execute the SQL scripts to create the MQTs.

Administration Tool Keyboard Shortcuts

This appendix provides keyboard shortcut information for the Administration Tool. The following sections list Administration Tool menu items and their corresponding keyboard shortcuts, keyboard shortcuts for navigating dialogs, and Physical Diagram and Business Model Diagram keyboard shortcuts.

This appendix contains the following topics:

- [Menu Keyboard Shortcuts](#)
- [Dialog Keyboard Shortcuts](#)
- [Physical Diagram and Business Model Diagram Keyboard Shortcuts](#)

Menu Keyboard Shortcuts

The following sections describe keyboard shortcuts for Administration Tool menu options.

File Menu Shortcuts

- The shortcut for New is Ctrl + N.
- The shortcut for Open, and then Offline is Ctrl + F.
- The shortcut for Open, and then Online is Ctrl + L.
- The shortcut for Save is Ctrl + S.
- The shortcut for Check Global Consistency is Ctrl + K.

Edit Menu Shortcuts

- The shortcut for Cut is Ctrl + X.
- The shortcut for Copy is Ctrl + C.
- The shortcut for Paste is Ctrl + V.
- The shortcut for Delete is Delete.

View Menu Shortcut

- The shortcut for Refresh is F5.

Tools Menu Shortcuts

- The shortcut for Show Consistency Checker is Ctrl + E.
- The shortcut for Query Repository is Ctrl + Q.

General Menu Shortcuts

Table G–1 lists the general keyboard shortcuts available in the Administration Tool menus. Note that you can use the Window menu options to change the focus from the menus to the navigation panes.

Table G–1 *Menu Keyboard Shortcuts*

Action	Keyboard Shortcut
Quit the application	Alt+ F4
Move cursor to the menu option	Alt + Underlined letter
Open application's control menu	Alt+ Spacebar
View the shortcut menu for the selected item	Shift + F10
Move through the menu bar	Left arrow key Right arrow key
Open a menu option	Down arrow key
Move through a menu list	Up arrow Down arrow
Close the current menu	Esc
Select or deselect items in a check box or list	Spacebar
Make noncontiguous selections	Ctrl + Up arrow + Spacebar

Dialog Keyboard Shortcuts

Table G–2 lists the keyboard shortcuts available in Administration Tool dialogs.

Table G–2 *Dialog Keyboard Shortcuts*

Action	Keyboard Shortcut
Move forward through options	Tab
Move backward through options	Shift + Tab
Select or deselect an item in a list	Shift + Up arrow Shift + Down arrow
Close the current dialog	Esc
Go to the top of a list	Home
Go to the bottom of a list	End
Refresh	F5
For dialogs with up arrow buttons: Move selected item up in the list	Alt + Up arrow Note: Select a list item before using this shortcut.

Table G-2 (Cont.) Dialog Keyboard Shortcuts

Action	Keyboard Shortcut
For dialogs with down arrow buttons: Move selected item down in the list	Alt + Down arrow Note: Select a list item before using this shortcut.
For dialogs with plus (add) buttons: Insert item from list	Alt + Insert
For dialogs with x (delete) buttons: Delete item from list	Alt + Delete
For dialogs with pencil (edit) buttons: Edit item from list	Alt + Enter
Browse dialog: Move focus between trees located in left pane	F5, F6, Shift + Tab, Tab
When a table row has a child row (grid): Expand the child row from the cell displaying the plus icon This situation occurs in the Define Merge Strategy page of the Merge Repository Wizard.	Spacebar Note: Move the focus to the cell displaying the plus icon before using this shortcut.
When a table row has a check box: Select or deselect the check box This situation occurs in the Define Merge Strategy page of the Merge Repository Wizard.	Spacebar Note: Move the focus to the cell displaying the check box before using this shortcut.

Physical Diagram and Business Model Diagram Keyboard Shortcuts

[Table G-3](#) lists the keyboard shortcuts available in the Physical and Business Model Diagrams. Note that the Physical and Business Model Diagram toolbar options are also available from the Actions menu.

Table G-3 Diagram Keyboard Shortcuts

Action	Keyboard Shortcut
Pan around the diagram when no diagram objects are selected	Arrow keys
Select a diagram object: use the arrow keys to move an object under the pointer, then press the spacebar to select the object	Spacebar + Arrow keys
Open the property dialog for a selected diagram object	Enter
Cancel current operation	Esc

Table G-3 (Cont.) Diagram Keyboard Shortcuts

Action	Keyboard Shortcut
Resume default mode (Select) after using Pan or Marquee Zoom	Esc
Deselect an object	Use one of the following methods: Esc Press the spacebar when the mouse cursor is not over an object
Zoom in	+
Zoom out	-
Select the pan tool	P Note that you can also use the arrow keys to pan around the diagram.
Revert to auto-layout	S
Create a new join	J This shortcut selects the New Join option.
Create a new table	N This shortcut selects the New Table option. After using this shortcut, you can use the arrow keys and spacebar to pan around the diagram and open the Logical Table dialog for the new table.
Select the Marquee Zoom tool	Z This shortcut selects the Marquee Zoom tool.
Zoom to fit all objects in the current view	F
Show all tables in Expanded View, with columns visible	1
Show all tables in Collapsed View, with columns hidden and only the table name displayed	2

Glossary

This glossary defines terms for Oracle Business Intelligence Enterprise Edition. See also the Oracle Fusion Middleware Master Glossary for additional terms and definitions.

action

Provides functionality to navigate to related content or to invoke operations, functions or processes in external systems. You can include actions in analyses, dashboard pages, agents, scorecard objectives, scorecard initiatives, and KPIs.

See also [action link](#).

Action Framework

The Action Framework is a component of the Oracle BI EE architecture and includes a J2EE application called the Action Execution Service (AES) and actions-specific JavaScript functionality deployed as part of Oracle BI EE. The action framework also includes client-side functionality for creating actions and invoking certain action types directly from the browser.

action link

A link to an action that you have embedded in an analysis, dashboard page, scorecard objective, scorecard initiative, or KPI that, when clicked, runs an associated action.

See also [action](#).

ADF Business Intelligence Component

Provides the developer the ability to include Oracle BI Presentation Catalog objects in ADF Applications. This component uses a SOAP connection to access the catalog.

Administration Server

Part of the WebLogic server domain and runs the processes that manage Oracle Business Intelligence components. The Administration Server includes the Oracle WebLogic Server Administration Console, Oracle Fusion Middleware Control, and JMX MBeans. For a Simple Install type, the Administration Server also includes Java components for Oracle Business Intelligence such as Oracle BI Publisher and Oracle Real-Time Decisions.

See also [Fusion Middleware Control](#), [Java components](#) and [Managed Server](#).

Administration Tool

See [Oracle BI Administration Tool](#).

agent

Enables you to automate your business processes. You can use agents to provide event-driven alerting, scheduled content publishing, and conditional event-driven action execution.

Agents can dynamically detect information-based problems and opportunities, determine the appropriate individuals to notify, and deliver information to them through a wide range of devices (e-mail, phones, and so on).

aggregate persistence

A feature that automates the creation and loading of aggregate tables and their corresponding Oracle Business Intelligence metadata mappings to enable aggregate navigation.

aggregate table

A table that stores precomputed results from measures that have been aggregated over a set of dimensional attributes. Each aggregate table column contains data at a given set of levels. For example, a monthly sales table might contain a precomputed sum of the revenue for each product in each store during each month. Using aggregate tables optimizes performance.

aggregation rule

In an Oracle BI repository, a rule applied to a logical column or physical cube column that specifies a particular aggregation function to be applied to the column data, such as SUM.

In Presentation Services, users can see the rules that have been applied in the repository. Users can also change the default aggregation rules for measure columns.

alias table

A physical table that references a different physical table as its source. You can use alias tables to set up multiple tables, each with different keys, names, or joins, when a single physical table must serve in different roles. Because alias table names are included in physical SQL queries, you can also use alias tables to provide meaningful table names, making the SQL statements easier to read.

analysis

A query that a user creates on the Criteria tab in Presentation Services. An analysis can optionally contain one or more filters or selection steps to restrict the results.

See also [filter](#) and [selection step](#).

analysis criteria

Consists of the columns, filters, and selection steps that you specify for an analysis.

See also [analysis](#).

analysis prompt

A prompt that is added to an analysis. When the user selects a prompt value, that value then determines the content that displays in the analysis that contains the prompt, only.

See [dashboard prompt](#) and [prompt](#).

attribute

The details of a dimension in an Oracle BI repository. Attributes usually appear as columns of a dimension table.

attribute column

In Presentation Services, a column that holds a flat list of values that are also known as members. No hierarchical relationship exists between these members, as is the case for members of a hierarchical column. Examples include ProductID or City.

See [hierarchical column](#).

BI Composer

BI Composer is a simple-to-use wizard that enables you to quickly and easily create, edit, or view analyses without the complexities of the Analysis editor.

BI domain

Contains configurable system components (the coreapplication) and Java components (the WebLogic server domain), and includes the Web-based management tools and applications that use resources.

A BI domain can be a set of middleware homes spread across one or more physical servers.

See also [BI instance](#).

BI instance

Refers to the system components (coreapplication) of a BI domain

See also [BI domain](#).

BI object

A piece of business intelligence content that is created with Presentation Services and saved to the Oracle BI Presentation Catalog. Examples of BI objects include analyses, dashboards, dashboard pages, scorecards, and KPIs.

BI Search

A search tool that resides outside of Presentation Services. BI Search is available from the Home Page after the administrator adds a link to the BI Search URL. BI Search provides a mechanism for searching for objects in the Oracle BI Presentation Catalog that is similar to a full-text search engine.

bookmark link

Captures the path to a dashboard page and all aspects of the page state.

See [prompted link](#).

bridge table

A table that enables you to resolve many-to-many relationships between two other tables.

briefing book

See [Oracle BI Briefing Books](#).

business model

Contains the business model definitions and the mappings from logical to physical tables. Business models are always dimensional, unlike objects in the Physical layer,

which reflect the organization of the data sources. Each business model contains logical tables, columns, and joins.

Business Model and Mapping layer

A layer of the Oracle BI repository that defines the business, or logical, model of the data and specifies the mapping between the business model and the Physical layer schemas. This layer can contain one or more business models.

The Business Model and Mapping layer determines the analytic behavior that is seen by users, and defines the superset of objects available to users. It also hides the complexity of the source data models.

business owner

The person responsible for managing and improving the business value and performance of a KPI or scorecard object, such as an objective, cause & effect map, and so on.

catalog

See [Oracle BI Presentation Catalog](#).

cause & effect map

A component of a scorecard that lets you illustrate the cause and effect relationships of an objective or KPI.

See also [Oracle Scorecard and Strategy Management](#).

chronological key

A column in a time dimension that identifies the chronological order of the members within a dimension level. The key must be unique at its level.

Cluster Controller

A process that serves as the first point of contact for new requests from Presentation Services and other clients. The Cluster Controller determines which Oracle BI Server in the cluster to direct the request to based on Oracle BI Server availability and load. It monitors the operation of servers in the cluster, including the Oracle BI Scheduler instances. The Cluster Controller is deployed in active-passive configuration.

column

In an Oracle BI repository, columns can be physical columns, logical columns, or presentation columns.

In Presentation Services, indicates the pieces of data that an analysis returns. Together with filters and selection steps, columns determine what analyses contain. Columns also have names that indicate the types of information that they contain, such as Account and Contact.

See also [analysis](#), [attribute column](#), [hierarchical column](#), and [measure column](#).

column filter

See [filter](#).

column prompt

A type of filter that enables you to build specific value prompts on a data column to either exist alone on the dashboard or analysis or to expand or refine existing dashboard and analysis filters.

See also [prompt](#).

complex join

A join in the Physical layer of an Oracle BI repository that uses an expression other than equals.

condition

Objects that return a single Boolean value based on the evaluation of an analysis or of a key performance indicator (KPI). You use conditions to determine whether agents deliver their content and execute their actions, whether actions links are displayed in dashboard pages, or whether sections and their content are displayed in dashboard pages.

See also [action](#), [action link](#), [agent](#) and [key performance indicator \(KPI\)](#).

connection pool

An object in the Physical layer of an Oracle BI repository that contains the connection information for a data source.

See also [Physical layer](#).

content designer

The user who creates business intelligence objects such as analyses, dashboards, and scorecards.

contextual event action

A predelivered action that uses the Action Framework to pass content from the business intelligence object to another region on an ADF page.

See also [action](#), [Action Framework](#), and [action link](#).

criteria

See [analysis criteria](#).

cube

An OLAP (online analytical processing) data structure that lets data be analyzed more quickly and with greater flexibility than structures in relational databases. Cubes are made up of measures and organized by dimensions. Cubes in multidimensional data sources roughly correspond to star schemas in relational database models.

currency prompt

A prompt that enables the user to change the currency type that displays in the currency columns on an analysis or dashboard.

See also [prompt](#).

custom view

A component of a scorecard that lets you show a customized view of your business and strategy data.

See also [Oracle Scorecard and Strategy Management](#).

dashboard

An object that provides personalized views of corporate and external information. A dashboard consists of one or more pages. Pages can display anything that you can

access or open with a Web browser, such as results of analyses, images, alerts from agents, and so on.

dashboard prompt

A prompt that is added to the dashboard. When the user selects a prompt value, that value then determines the content that displays in all analyses that are included on the dashboard.

See [analysis prompt](#) and [prompt](#).

Dashboard URL

Used for incorporating or referencing the content of a specific dashboard in external portals or applications. It has several forms and optional arguments that you can use to control its behavior.

data source name (DSN)

A data structure that contains the information about a specific database, typically used by an ODBC driver to connect to the database. The DSN contains information such as the name, directory, and driver of the database.

Connection pool objects in the Physical layer of the Oracle BI repository contain DSN information for individual data sources.

database hint

Instructions placed within a SQL statement that tell the database query optimizer the most efficient way to execute the statement. Hints override the optimizer's execution plan, so you can use hints to improve performance by forcing the optimizer to use a more efficient plan. Hints are supported only for Oracle Database data sources.

dimension

A hierarchical organization of logical columns (attributes). One or more logical dimension tables might be associated with at most one dimension.

A dimension might contain one or more (unnamed) hierarchies. There are two types of logical dimensions: dimensions with level-based hierarchies (structure hierarchies), and dimensions with parent-child hierarchies (value hierarchies).

A particular type of level-based dimension, called a time dimension, provides special functionality for modeling time series data.

See also [hierarchy](#).

dimension table

A logical table that contains columns used by a particular dimension. A dimension table cannot be a fact table.

See also [fact table](#).

driving table

A mechanism used to optimize the manner in which the Oracle BI Server processes multi-database joins when one table is very small (the driving table) and the other table is very large.

DSN

See [data source name \(DSN\)](#).

event polling table

Event polling tables (also called event tables) provide information to the Oracle BI Server about which physical tables have been updated. They are used to keep the query cache up-to-date. The Oracle BI Server cache system polls the event table, extracts the physical table information from the rows, and purges stale cache entries that reference those physical tables.

Essbase

A multidimensional database management system available from Oracle that provides a multidimensional database platform upon which to build business intelligence applications. Also referred to as Oracle's Hyperion Essbase.

fact table

In an Oracle BI repository, a logical table in the Business Model and Mapping layer that contains measures and has complex join relationships with dimension tables.

See also [dimension table](#).

filter

Criteria that are applied to attribute and measure columns to limit the results that are displayed when an analysis is run. For measure columns, filters are applied before the query is aggregated and affect the query and thus the resulting values.

See also [prompt](#) and [selection step](#).

foreign key

A column or a set of columns in one table that references the primary key columns in another table.

fragmentation content

The portion, or fragment, of the set of data specified in a logical table source when the logical table source does not contain the entire set of data at a given level.

Fragmentation content is defined by the logical columns that are entered in the **Fragmentation content** box in the Content tab of the Logical Table Source dialog box.

Fusion Middleware Control

Provides Web-based management tools that enable you to monitor and configure Fusion Middleware components.

global header

An object in the user interface for Oracle BI Presentation Services that contains links and options that enable the user to quickly begin a task or locate a specific object within the Oracle BI Presentation Catalog. The global header always displays in the Presentation Services user interface, thus enabling users to quickly access links and search the catalog without having to navigate to the Home Page or Catalog page.

Go URL

Used to incorporate specific business intelligence results into external portals or applications. The Go URL is used when you add a result to your favorites or add a link to a request to a dashboard or external Web site. It has several forms and optional arguments that you can use to control its behavior.

hierarchical column

In Presentation Services, a column that holds data values that are organized using both named levels and parent-child relationships. This column is displayed using a tree-like structure. Individual members are shown in an outline manner, with lower-level members rolling into higher-level members. For example, a specific day belongs to a particular month, which in turn is within a particular year. Examples include Time or Geography.

See also [attribute column](#).

hierarchy

In an Oracle BI repository, a system of levels in a logical dimension that are related to each other by one-to-many relationships. All hierarchies must have a common leaf level and a common root (all) level.

Hierarchies are not modeled as separate objects in the metadata. Instead, they are an implicit part of dimension objects.

See also [dimension](#), [logical level](#), and [presentation hierarchy](#).

hierarchy level

In Presentation Services, an object within a hierarchical column that either rolls up or is rolled up from other levels. Corresponds to a presentation level in an Oracle BI repository.

See also [presentation level](#).

home page

Provides an intuitive, task-based entry way into the functionality of Presentation Services. The Home page is divided into sections that enable you to quickly begin specific tasks, locate an object, or access technical documentation.

image prompt

A prompt that provides an image with different areas mapped to specific values. The user clicks an image area to select the prompt value that populates the analysis or dashboard.

See also [prompt](#).

initialization block

Used to initialize dynamic repository variables, system session variables, and nonsystem session variables. An initialization block contains the SQL statements that are executed to initialize or refresh the variables that are associated with that block.

initiative

Used in a scorecard, an initiative is a time-specific task or project that is necessary to achieve objectives. As such, you can use initiatives that support objectives as milestones as they reflect progress toward strategy targets.

See also [objective](#) and [Oracle Scorecard and Strategy Management](#).

Java components

Fusion Middleware Control components that are deployed as one or more Java EE applications (and a set of resources) and are managed by Node Manager.

See also [Node Manager](#).

key performance indicator (KPI)

A measurement that defines and tracks specific business goals and strategic objectives. KPIs often times roll up into larger organizational strategies that require monitoring, improvement, and evaluation. KPIs have measurable values that usually vary with time, have targets to determine a score and performance status, include dimensions to allow for more precise analysis, and can be compared over time for trending purposes and to identify performance patterns.

See also [Oracle Scorecard and Strategy Management](#).

KPI watchlist

A method of distributing KPIs to end users. A watchlist is a collection of KPIs that are built by adding the KPIs that are stored in the Oracle BI Presentation Catalog. After a KPI watchlist is built and saved, it is stored as a catalog object and can be added to dashboards and scorecards.

See also [key performance indicator \(KPI\)](#).

level

See [hierarchy level](#).

logical display folder

Folders used to organize objects in the Business Model and Mapping layer. They have no metadata meaning.

logical join

Joins that express relationships between logical tables. Logical joins are *conceptual*, rather than physical, joins. In other words, they do not join to particular keys or columns. A single logical join can correspond to many possible physical joins.

logical layer

See [Business Model and Mapping layer](#).

logical level

In an Oracle BI repository, a component of a level-based hierarchy that either rolls up or is rolled up from other levels.

Parent-child hierarchies have implicit, inter-member levels between ancestors and descendants that are not exposed as logical level objects in the metadata. Although parent-child hierarchies also contain logical level objects, these levels are system generated and exist to enable aggregation across all members only.

See also [dimension](#) and [hierarchy](#).

Logical SQL

The SQL statements that are understood by the Oracle BI Server. The Oracle BI Server Logical SQL includes standard SQL, plus special functions (SQL extensions) like `AGO`, `TODATE`, `EVALUATE`, and others.

Clients like Presentation Services send Logical SQL to the Oracle BI Server when a user makes a request. In addition, Logical SQL is used in the Business Model and Mapping layer to enable heterogeneous database access and portability. The Oracle BI Server transforms Logical SQL into physical SQL that can be understood by source databases.

logical table

A table object in the Business Model and Mapping layer of an Oracle BI repository. A single logical table can map to one or more physical tables. Logical tables can be either fact tables or dimension tables.

See also [dimension table](#) and [fact table](#).

logical table source

Objects in the Business Model and Mapping layer of an Oracle BI repository that define the mappings from a single logical table to one or more physical tables. The physical to logical mapping can also be used to specify transformations that occur between the Physical layer and the Business Model and Mapping layer, and to enable aggregate navigation and fragmentation.

Managed Server

An individual J2EE application container (JMX MBean container). It provides local management functions on individual hosts for Java components and system components contained within the local middleware home, and refers to the Administration Server for all of its configuration and deployment information.

See also [Administration Server](#) and [Fusion Middleware Control](#).

measure column

A column that can change for each record and can be added up or aggregated. Typical measures are sales dollars and quantity ordered. Measures are calculated from data sources at query time.

Measure columns are displayed in the Oracle BI repository, usually in fact tables, or in Presentation Services.

metadata

Data about data. Metadata objects include the descriptions of schemas (such as tables, columns, data types, primary keys, foreign keys, and so on) and logical constructs (like fact tables, dimensions, and logical table source mappings).

The Oracle BI repository is made up of the metadata used by the Oracle BI Server to process queries.

metadata dictionary

A static set of XML documents that describe metadata objects, such as a column, including its properties and relationships with other metadata objects. A metadata dictionary can help users obtain more information about metrics or attributes for repository objects.

mission statement

A statement in a scorecard that specifies the key business goals and priorities that are required to achieve your vision.

See also [Oracle Scorecard and Strategy Management](#) and [vision statement](#).

multi-database join

A join between two tables in an Oracle BI repository, where each table resides in a different database.

Node Manager

A daemon process that provides remote server start, stop, and restart capabilities when Java processes become unresponsive or terminate unexpectedly.

See also [Java components](#).

OCI

See [Oracle Call Interface \(OCI\)](#).

ODBC

See [Open Database Connectivity \(ODBC\)](#).

object properties

Information about an object and attributes that the owner can assign to an object. Examples of properties include name, description, date stamps, read-only access, and do not index flag.

See also [permissions](#).

objective

A required or desired outcome in a scorecard that forms your corporate strategy.

See also [initiative](#) and [Oracle Scorecard and Strategy Management](#).

offline mode

In the Oracle BI Administration Tool, a mode where a repository builder can edit a repository that is not loaded into the Oracle BI Server.

online mode

In the Oracle BI Administration Tool, a mode where a repository builder can edit a repository while it is available for query operations. Online mode also allows user session monitoring for users connected to the subject areas in the repository.

opaque view

A Physical layer table that consists of a `SELECT` statement. In the Oracle BI repository, opaque views appear as view tables in the physical databases, but the view does not actually exist.

Open Database Connectivity (ODBC)

A standard interface used to access data in both relational and nonrelational databases. Database applications can use ODBC to access data stored in different types of database management systems, even if each database uses a different data storage format and programming interface.

OPMN

See [Oracle Process Manager and Notification Server \(OPMN\)](#).

Oracle BI Administration Tool

A Windows application that is used to create and edit Oracle BI repositories. The Administration Tool provides a graphical representation of the three parts of a repository: the Physical layer, the Business Model and Mapping layer, and the Presentation layer.

Oracle BI Briefing Books

A collection of static or updatable snapshots of dashboard pages, individual analyses, and BI Publisher reports. You can download briefing books in PDF or MHTML format for printing and viewing. You also can update, schedule, and deliver briefing books using agents.

Oracle BI JavaHost

A service that gives Presentation Services the ability to use functionality that is provided in Java libraries to support components such as graphs. The services are provided based on a request-response model.

Oracle BI Logical SQL View Object

Provides the developer the ability to create a Logical SQL statement to access the Oracle BI Server and fetch business intelligence data and bind it to native ADF components for inclusion on an ADF page. This view object uses a BI JDBC connection to the Oracle BI Server.

Oracle BI Presentation Catalog

Stores business intelligence objects, such as analyses and dashboards, and provides an interface where users create, access, and manage objects, and perform specific object-based tasks (for example, export, print, and edit). The catalog is organized into folders that are either shared or personal.

Oracle BI Presentation Services

Provides the framework and interface for the presentation of business intelligence data to Web clients. It maintains a Presentation Catalog service on the file system for the customization of this presentation framework. It is a standalone process and communicates with the Oracle BI Server using ODBC over TCP/IP. It consists of components that are known as Answers, Delivers, and Interactive Dashboards.

See also [ODBC](#); [Oracle BI Server](#); [Oracle BI Presentation Catalog](#); [Oracle BI Presentation Services server](#).

Oracle BI Presentation Services server

The Oracle BI Web server that exchanges information and data with the Oracle BI Server.

Oracle BI Publisher

A J2EE application that provides enterprise-wide publishing services in Oracle Business Intelligence. It generates highly formatted, pixel-perfect reports.

See also [report](#).

Oracle BI Publisher report

See [report](#).

Oracle BI repository

A file that stores Oracle Business Intelligence metadata. The metadata defines logical schemas, physical schemas, physical-to-logical mappings, aggregate table navigation, and other constructs. The repository file has an extension of .rpd. Oracle BI repositories can be edited using the Oracle BI Administration Tool.

See also [metadata](#) and [Oracle BI Administration Tool](#).

Oracle BI Scheduler

An extensible scheduling application for scheduling results to be delivered to users at specified times. It is the engine behind the Oracle BI Delivers feature.

See also [results](#).

Oracle BI Server

A standalone process that maintains the logical data model that it provides to Presentation Services and other clients through ODBC. Metadata is maintained for the data model in a local proprietary file called the repository file. The Oracle BI Server processes user requests and queries underlying data sources.

Oracle BI Server XML API

Provides utilities to create a generic, XML-based representation of the Oracle BI repository metadata. You can use this XML file version of the repository to programmatically modify the metadata. The Oracle BI Server XML API objects correspond to metadata repository objects in an RPD file. These objects differ from XML objects in the Oracle BI Presentation Catalog.

Oracle Business Intelligence Mobile

Oracle Business Intelligence Mobile allows you to view Oracle BI EE content on supported mobile devices such as the Apple iPhone and Apple iPad.

Using Oracle Business Intelligence Mobile, you can view and analyze BI content such as analyses and dashboards, BI Publisher content, scorecard content, and content delivered by agents.

Oracle Business Intelligence Web Services

See [Oracle Business Intelligence Session-Based Web Services](#) and [Oracle Business Intelligence Web Services for SOA](#).

Oracle Business Intelligence Session-Based Web Services

An API that implements SOAP. These Web services are designed for programmatic use, where a developer uses one Web service to invoke many different business intelligence objects. These Web services provide functionality on a wide range of Presentation Services operations. These Web services enable the developer to extract results from Oracle BI Presentation Services and deliver them to external applications, perform Presentation Services management functions, and execute Oracle Business Intelligence alerts (known as Intelligent Agents).

See also [Oracle Business Intelligence Web Services for SOA](#).

Oracle Business Intelligence Web Services for SOA

Contains three Web services, ExecuteAgent, ExecuteAnalysis, and ExecuteCondition, which are hosted by the bimiddleware J2EE application. These web services are designed to enable developers to use third-party Web services clients (for example, Oracle SOA Suite) to browse for and include business intelligence objects in service oriented architecture components.

See also [Oracle Business Intelligence Session-Based Web Services](#).

Oracle Call Interface (OCI)

A connection interface that the Oracle BI Server can use to connect to Oracle Database data sources. You should always use OCI when importing metadata from or connecting to an Oracle Database.

Oracle OLAP

Oracle Database has an OLAP Option that provides an embedded, full-featured online analytical processing server.

Oracle Business Intelligence supports Oracle OLAP as a data source. When you import metadata from an Oracle OLAP source, the Oracle OLAP objects appear in the Physical layer of the Administration Tool. Oracle OLAP objects include Analytic Workspaces, which are containers for storing related cubes.

Oracle Process Manager and Notification Server (OPMN)

A process management tool that manages all system components (server processes), and supports both local and distributed process management, automatic process recycling and the communication of process state (up, down, starting, stopping). OPMN detects process unavailability and automatically restarts processes).

See also [system components](#).

Oracle Scorecard and Strategy Management

A performance management tool that lets you describe and communicate your business strategy. You can drive and assess your corporate strategy and performance from the top of your organization down, or from the bottom up.

Oracle Technology Network (OTN)

A repository of technical information about Oracle's products where you can search for articles, participate in discussions, ask the user community technical questions, and search for and download Oracle products and documentation.

parent-child hierarchy

A hierarchy of members that all have the same type. All the dimension members of a parent-child hierarchy occur in a single data source. In a parent-child hierarchy, the inter-member relationships are parent-child relationships between dimension members.

See also [dimension](#).

parent-child relationship table

A table with values that explicitly define the inter-member relationships in a parent-child hierarchy. Also called a closure table.

pass-through calculation

A calculation that is not computed by the Oracle BI Server but instead is passed to another data source. Enables advanced users to leverage data source features and functions without the need to modify the Oracle BI repository.

permissions

Specify which users can access an object, and limit how users can interact with an object. Examples of permissions include write, delete, and change permissions.

See [object properties](#).

perspective

A category in your organization with which to associate initiatives, objectives, and KPIs in a scorecard. A perspective can represent a key stakeholder (such as a customer, employee, or shareholder/financial) or a key competency area (such as time, cost, or quality).

See also [initiative](#), [key performance indicator \(KPI\)](#), [objective](#), and [Oracle Scorecard and Strategy Management](#).

physical catalog

An object in the Physical layer of a repository that groups different schemas. A catalog contains all the schemas (metadata) for a database object.

physical display folder

Folders that organize objects in the Physical layer of an Oracle BI repository. They have no metadata meaning.

physical join

Joins between tables in the Physical layer of an Oracle BI repository.

Physical layer

A layer of the Oracle BI repository that contains objects that represent physical data constructs from back-end data sources. The Physical layer defines the objects and relationships available for writing physical queries. This layer encapsulates source dependencies to enable portability and federation.

physical schema

An object in the Physical layer of an Oracle BI repository that represents a schema from a back-end database.

physical table

An object in the Physical layer of an Oracle BI repository, usually corresponding to a table that exists in a physical database.

See also [Physical layer](#).

presentation hierarchy

An object in the Presentation layer of an Oracle BI repository that provides an explicit way to expose the multidimensional model in Presentation Services and other clients. Presentation hierarchies expose analytic functionality such as member selection, custom member groups, and asymmetric queries. Users can create hierarchy-based queries using presentation hierarchies.

In Presentation Services, presentation hierarchies are displayed as hierarchical columns.

See also [hierarchical column](#) and [presentation level](#).

Presentation layer

Provides a way to present customized, secure, role-based views of a business model to users. It adds a level of abstraction over the Business Model and Mapping layer in the Oracle BI repository. The Presentation layer provides the view of the data seen by users who build analyses in Presentation Services and other client tools and applications.

See also [Business Model and Mapping layer](#).

presentation level

In the Oracle BI repository, a component of a presentation hierarchy that either rolls up or is rolled up from other levels. Presentation levels are displayed as levels within hierarchical columns in Presentation Services.

See also [hierarchy level](#) and [presentation hierarchy](#).

Presentation Services

See [Oracle BI Presentation Services](#).

Presentation Services server

See [Oracle BI Presentation Services server](#).

presentation table

An object in the Presentation layer of an Oracle BI repository that is used to organize columns into categories that make sense to the user community. A presentation table can contain columns from one or more logical tables. The names and object properties of the presentation tables are independent of the logical table properties.

primary key

A column (or set of columns) where each value is unique and identifies a single row of a table.

process instance

A unique process on an individual workstation that is associated with a BI instance.

See also [BI instance](#).

prompt

A type of filter that enables the content designer to build and specify data values or the end user to choose specific data values to provide a result sets for an individual analysis or multiple analyses included on a dashboard or dashboard page. A prompt expands or refines existing dashboard and analysis filters.

The types of prompts are column prompts, currency prompts, image prompts, and variable prompts.

See also [column prompt](#), [currency prompt](#), [filter](#), [image prompt](#), and [variable prompt](#).

prompted link

Captures the path to a dashboard page and a simplified presentation of the dashboard prompt.

See [bookmark link](#).

query

Contains the underlying SQL statements that are issued to the Oracle BI Server. You do not have to know a query language to use Oracle Business Intelligence.

query cache

A facility to store query results for use by other queries.

ragged hierarchy

See [unbalanced hierarchy](#).

report

The response returned to the user from the execution of a query created using Oracle BI Publisher. Reports can be formatted, presented on a dashboard page, saved in the Oracle BI Presentation Catalog, and shared with other users.

See also [analysis](#).

repository

See [Oracle BI repository](#).

repository variable

See [variable](#).

results

The output returned from the Oracle BI Server for an analysis.

See also [analysis](#).

scorecard

See [Oracle Scorecard and Strategy Management](#).

selection step

A choice of values that is applied after the query is aggregated that affects only the members displayed, not the resulting aggregate values. Along with filters, selection steps restrict the results for an analysis.

See also [analysis](#) and [filter](#).

session variable

See [variable](#).

skip-level hierarchy

A hierarchy where some members do not have a value for a particular ancestor level. For example, in the United States, the city of Washington in the District of Columbia does not belong to a state. The expectation is that users can still navigate from the country level (United States) to Washington and below without the need for a state.

See also [hierarchy](#).

snowflake schema

A dimensional schema where one or more of the dimensions are partially or completely normalized.

SQL

See [structured query language \(SQL\)](#).

star schema

A relational schema that allows dimensional analysis of historical information. Star schemas have one-to-many relationships between the logical dimension tables and the logical fact table. Each star consists of a single fact table joined to a set of denormalized dimension tables.

strategy map

A component of a scorecard that shows how the objectives that have been defined for a scorecard and the KPIs that measure their progress are aligned by perspectives. It also shows cause and effect relationships.

See also [Oracle Scorecard and Strategy Management](#).

strategy tree

A component of a scorecard that shows an objective and its supporting child objectives and KPIs hierarchically in a tree diagram.

See also [Oracle Scorecard and Strategy Management](#).

structured query language (SQL)

A standard programming language for querying and modifying data. Oracle Business Intelligence supports standard SQL-92 with several value-added proprietary extensions.

See also [Logical SQL](#).

subject area

In an Oracle BI repository, an object in the Presentation layer that organizes and presents data about a business model. It is the highest-level object in the Presentation layer and represents the view of the data that users see in Presentation Services. Oracle BI repository subject areas contain presentation tables, presentation columns, and presentation hierarchies.

In Presentation Services, subject areas contain folders, measure columns, attribute columns, hierarchical columns, and levels.

system components

Server processes (not Java applications) that are managed by the Oracle Process Manager and Notification server (OPMN).

See also [Oracle Process Manager and Notification Server \(OPMN\)](#).

transformation

Work that is performed on data when moving from a database to another location (sometimes another database). Some transformations are typically performed on data when it is moved from a transaction system to a data warehouse system.

unbalanced hierarchy

A hierarchy where the leaves do not have the same depth. For example, an organization might choose to have data for the current month at the day level, data for the previous at the month level, and data for the previous five years at the quarter level.

See also [hierarchy](#).

value hierarchy

See [parent-child hierarchy](#).

variable

Objects in an Oracle BI repository that are used to streamline administrative tasks and dynamically modify metadata content to adjust to a changing data environment.

Variables are of the following types:

- Repository variables have a single value at any point in time. There are two types of repository variables: static and dynamic.
- Session variables are created and assigned a value when each user logs on. There are two types of session variables: system and nonsystem.

variable prompt

Enables the user to select a value specified in the variable prompt to display on the dashboard. A variable prompt is not dependent upon column data, but enables you to manipulate, for example add or multiply, the column data on an analysis.

See also [prompt](#).

virtual physical table

A physical table that is made from a stored procedure or a `SELECT` statement. Creating virtual tables can provide the Oracle BI Server and the underlying databases with the proper metadata to perform some advanced query requests.

vision statement

A short statement in a scorecard that describes what your organization wants to become sometime in the future. For example, it might be to become the most successful business in the South America Polypropylene Market.

See also [mission statement](#) and [Oracle Scorecard and Strategy Management](#).

WebLogic server domain

Contains Java components that are configured to participate in the servicing of SOAP, HTTP, and other forms of requests.

WebLogic Scripting Tool (WLST)

A command-line scripting interface that enables you to configure, manage, and persist changes to WebLogic Server instances and domains and to monitor and manage server run-time events.

XML API

See [Oracle BI Server XML API](#).

Numerics

503 Service Not Available errors, avoiding, 6-9

A

ABS math function, about, C-48

ACOS math function, about, C-48

Actions menu, about, 2-10

ADF Business Components

about, 5-1

application EAR file, deploying, 5-5

AppModule Config, 5-10, 6-20

AppModuleDefinition, 5-10, 6-20

enabling custom parameters, 5-11

importing metadata from, 5-9

JDBC data source for, 5-8

propagating labels and tooltips, 5-11

properties in connection pool, 6-19

setting up, 5-3

URL, 5-10, 6-20

working with, 5-1

Administration Tool

about, 2-1

Actions menu, 2-10

Edit menu, 2-8

File menu, 2-6

Help menu, 2-10

icons, 2-14

keyboard shortcuts, G-1

main window, 2-2

Manage menu, 2-8

opening, 2-2

setting preferences, 2-3

toolbar, 2-2

Tools menu, 2-9

version, 2-10

View menu, 2-8

Window menu, 2-10

Administration Tool utilities

Aggregate Persistence, 17-15

Externalize Strings, 17-10

Generate Metadata Dictionary, 17-14

Oracle BI Event Tables, 17-10

Remove Unused Physical Objects, 17-15

Rename Wizard, 17-11

Replace Column or Table in Logical Table

Sources, 17-7

Repository Documentation, 17-13

Update Physical Layer, 17-12

AGGREGATE AT aggregate function, about, C-25

aggregate functions

about, C-24, C-31

AGGREGATE AT, C-25

alternative syntax, C-13

AVG, C-25

AVGDISTINCT, C-26

BOTTOMN, C-26

computing aggregates of baseline columns, C-9

computing aggregates of measure columns, C-11

COUNT, C-26

COUNT (*), C-26

COUNTDISTINCT, C-26

display function reset behavior, C-12

FIRST, C-27

GROUPBYCOLUMN, C-27

GROUPBYLEVEL, C-28

LAST, C-28

MAX, C-28

MEDIAN, C-29

MIN, C-29

NTILE, C-29

PERCENTILE, C-29

RANK, C-30

rules for, C-9

STDDEV, C-30

STDDEV_POP, C-30

SUM, C-31

SUMDISTINCT, C-31

TOPN, C-31

using FILTER to compute conditional

aggregates, C-13

aggregate navigation

about aggregate table fragments, 10-15

about WHERE clause filters, 10-10

creating sources for each level of aggregated fact

data, 10-11

specifying aggregate levels for each source, 10-11

specifying fragmentation content, 10-12

aggregate persistence

about, 12-1

default prefix, 12-2

- identifying query candidates, 12-2
- running the specification, 12-10
- troubleshooting, 12-10
- using the Aggregate Persistence Wizard, 12-3
- writing script manually, 12-6
- Aggregate Persistence Wizard, about, 12-3
- aggregate queries, about, C-9
- aggregate table fragments
 - about, 10-15
 - about configuring a repository to use
 - fragments, 10-16
 - constructing physical joins for virtual table, 10-17
 - creating SQL virtual table content, 10-17
 - defining Physical layer tables with a SELECT Statement, 10-17
 - specifying aggregate table content, 10-16
- AGGREGATE_PREFIX configuration option, 12-2
- Aggregation content, group-by list, 10-8
- aggregation rules
 - default rules for measure columns, 8-15
 - dimension-specific, 8-16
 - Evaluate_Aggr, 8-15
 - for multidimensional data sources, 8-18
 - for Oracle OLAP, 7-31
- AGO time series function, about, C-36
- alias tables
 - about, 7-8
 - creating, 7-12
 - for Essbase, 7-24
- aliases
 - Alias tab, using, 11-14
 - merging, D-7
- Allow direct database requests by default
 - about, 6-3
 - recommendations for enabling, 6-3
- Allow first Connection Pool for Init Blocks
 - option, 2-5
- Allow import from repository option, 2-4
- Allow logical foreign key join creation option, 2-4
- Allow populate queries by default, 6-3
- alternate hierarchies, about, 7-25
- Analytic Workspace objects, working with, 7-31
- Ancestor Reference column type, 7-21
- annotations, enabling for Oracle Scorecard and Strategy Management, 14-1
- application EAR file, for ADF Business Components, 5-5
- application module, 5-2
- application roles
 - about granting privileges for, 13-10
 - managing, 1-16
 - managing in the default policy store, 13-2
 - permission inheritance, 13-10
 - setting up placeholder roles, 13-15
- AppModule Config, for ADF Business Components, 5-10, 6-20
- AppModule Definition, for ADF Business Components, 5-10, 6-20
- ARBORPATH, setting, 4-6, 15-13
- architecture

- multiuser development, A-4
- Oracle BI Server, 1-1
- repository, 1-3
- ASCII string function, about, C-39
- ASIN math function, about, C-48
- ATAN math function, about, C-48
- ATAN2 math function, about, C-48
- Attribute column type, 7-21
- auditing, for data access security, 13-1
- authentication
 - cache, 2-6
 - managing, 13-2
- Auto Layout option, for Diagrams, 2-10
- AVG aggregate function, about, C-25
- AVGDISTINCT aggregate function, about, C-26

B

- Based on dimensions option for measures, 8-16
- baseline column
 - behavior with aggregate functions, C-9
 - example, C-9
- Between SQL logical operator, about, C-14
- bi-init, 2-2, 2-16, 2-24, 3-11, 4-7, 14-3, 16-5, 16-20, A-6, E-2
- biserverxmlcli, 2-16
- biserverxmlexec, 2-16
- biserverxmlgen, 2-16
- BISystem user, 2-5
- BIT_LENGTH string function, about, C-40
- blank repository, creating, 16-15
- BOTTOMN aggregate function, about, C-26
- branching, in multiuser development, 3-17, A-16
- bridge tables
 - about, 8-23
 - modeling in a single dimension, 8-24
 - modeling in separate dimensions, 8-25
 - modeling in the Physical layer, 8-24
 - weight factor for, 8-23
- Browse dialog
 - about using, 2-13
 - querying for an object, 2-14
 - selecting an object in, 2-14
 - synchronizing an object in the query results list
 - with the tree control list, 2-14
- buffer size, setting up, 6-22
- bulk creation of metrics, 17-15
- Business Model and Mapping layer
 - about, 1-3, 8-1
 - associating measures with levels in a dimension, 8-19
 - automatically creating objects for, 8-2
 - creating logical columns, 8-11
 - creating logical table source content
 - definitions, 10-8
 - defining physical to logical mapping, 10-4
 - managing logical table sources, 10-1
 - setting up display folders in, 8-22
 - tips for, 1-12
 - working with logical tables, 8-4

- Business Model Diagram
 - about using, 2-10
 - about using to create joins, 8-7
 - displaying, 8-3
 - displaying qualified names in, 2-3
 - working with, 8-3
- business models
 - creating, 8-2
 - disabling, 8-2
 - duplicating with subject area, 8-3
 - planning, 1-4
 - understanding, 1-4

C

- Cache Manager
 - about, 2-9
 - column order, 2-6
- Cache menu option, 2-9
- Cache refresh interval option (authentication cache), 2-6
- caching
 - authentication cache, 2-6
 - managing, 1-16
 - settings for physical tables, 7-11
- calculated items, creating, 10-4, 10-5
- CALCULATEDMEMBER
 - examples, C-22
 - rules, C-20
 - syntax, C-19
- Calculation Wizard
 - about, 17-15
 - setting up, 2-4
 - using, 17-15
- calculations, for level-based measures, 9-7
- calendar date/time functions
 - about, C-53
 - CURRENT_DATE, C-54
 - CURRENT_TIME, C-54
 - CURRENT_TIMESTAMP, C-54
 - DAY_OF_QUARTER, C-55
 - DAYNAME, C-55
 - DAYOFMONTH, C-55
 - DAYOFWEEK, C-55
 - DAYOFYEAR, C-56
 - HOUR, C-56
 - MINUTE, C-56
 - MONTH, C-56
 - MONTH_OF_QUARTER, C-56
 - MONTHNAME, C-57
 - NOW, C-57
 - QUARTER_OF_YEAR, C-57
 - SECOND, C-57
 - TIMESTAMPADD, C-57
 - TIMESTAMPDIFF, C-59
 - WEEK_OF_QUARTER, C-60
 - WEEK_OF_YEAR, C-60
 - YEAR, C-61
- call interface, for connection pools, 6-8
- CASE (If) conditional expression, about, C-16

- CASE (Switch) conditional expression, about, C-15
- CAST conversion function, about, C-61, C-62
- catalogs, creating, 7-4
- CEILING math function, about, C-49
- certification information, 1-16
- CHAR string function, about, C-40
- CHAR_LENGTH string function, about, C-40
- character literals, about, C-17
- Check out objects automatically option, 2-4
- checking global consistency, 2-21
- checking in changes, 2-20
- checking out objects, 2-20
- chronological keys
 - about, 9-3
 - selecting and sorting, 9-10
 - setting for each level, 9-26
- circular joins, about and eliminating, 1-11
- client applications, enabling connectivity for, 14-5
- Cluster Manager, 2-9
- Cluster menu option, 2-9
- Collapse All option, for Diagrams, 2-11
- column mapping, logical to physical, 10-4
- column types, for Essbase, 7-21
- command-line utilities, 2-15
- Common Enterprise Information Model, about, 1-4
- Compare Mode, turning off, 16-3
- Compare with Original menu option, 3-12
- comparing repositories, 16-1
- complex joins
 - See physical complex joins
- CONCAT string function, about, C-40
- conditional expressions
 - about, C-15
 - CASE (If), about and syntax, C-16
 - CASE (Switch), about and syntax, C-15
- conformed dimensions, 1-6
- connection pools
 - about, 1-11, 6-5
 - automating changes for, 6-7
 - call interface, 6-8
 - creating connection scripts, 6-14
 - creating or changing, 6-6
 - data source name, 6-10
 - for initialization blocks, 6-6
 - for standby database configuration, 4-36
 - maximum number of connections, 6-9
 - permissions for, 6-8
 - persist connection pool property, 6-21
 - properties for ADF Business Components, 6-19
 - require fully qualified table names, 6-10
 - shared logon, 6-11
 - Unicode database type option, 6-19
 - URL, 6-13
 - using net service name, xxxi
 - write-back properties, 6-17
 - XML properties, 6-16
- connection scripts, in connection pool, 6-14
- consistency check
 - checking repository for consistency, 2-23
 - checking single object for consistency, 2-23

- new validation rules, xxiv, xxxi
- Consistency Check Manager
 - about checking repository or objects for consistency, 2-21
 - about messages, 2-22
 - about passing consistency check, 2-21
 - checking repository for consistency, 2-23
 - checking single object for consistency, 2-23
 - copying inconsistency messages, 2-23
 - correcting inconsistencies, 2-23
- context for dimension-only queries, specifying, 11-4, 11-6
- conversion functions
 - about, C-61
 - CAST, C-61
 - CHOOSE, C-62
 - IFNULL, C-62
 - INDEXCOL, C-62
 - TO_DATETIME, C-64
 - VALUEOF, C-65
- COS math function, about, C-49
- COT math function, about, C-49
- COUNT aggregate function, about, C-26
- COUNT(*) aggregate function, about, C-26
- COUNTDISTINCT aggregate function, about, C-26
- create aggregates specification, writing
 - manually, 12-6
- Create New Repository Wizard, 4-2
- create view function, about, 7-39
- create/prepare aggregates syntax, 12-9
- creating calculated items, 10-4, 10-5
- credential store, managing system credentials
 - in, 13-2
- cube variables, working with, 7-19
- cubes
 - about, 1-9, 7-15
 - viewing members, 7-20
 - working with, 7-1
- current repository, about, 16-8
- CURRENT_DATE calendar date/time function,
 - about, C-54
- CURRENT_TIME calendar date/time function,
 - about, C-54
- CURRENT_TIMESTAMP calendar date/time function, about, C-54
- custom authenticators, managing, 13-3
- custom properties, enabling for ADF Business Components, 5-11

D

- data access security
 - about, 13-1
 - about applying in offline mode, 13-14
 - auditing, 13-1
- data filters, setting up, 13-3
- Data is dense option, 8-16
- data modeling
 - identifying data source content, 1-8
 - identifying dimensions, 1-7

- identifying logical dimension tables, 1-6
- identifying logical fact tables, 1-5
- identifying lookup tables, 1-8
- objectives of, 1-4
- planning the business model, 1-4
- data source name, in connection pool, 6-10
- data sources
 - connecting on Linux and UNIX using native ODBC drivers, 15-10
 - DataDirect Connect drivers, 15-6
 - DB2 Connect, 15-13
 - Essbase setup, 4-6, 15-12
 - Hyperion Financial Management setup, 4-7
 - Microsoft Analysis Services, 4-16
 - ODBC DSN setup, 4-3
 - Oracle Database setup, 4-4
 - Oracle OLAP setup, 4-4
 - Oracle RPAS setup, 4-10, 15-11
 - preconfiguration tasks, 4-2
 - SAP/BW setup, 4-10
 - setup on Linux and UNIX, 15-1
 - supported by Oracle Business Intelligence, 1-16
 - Teradata setup, 4-4, 15-10
 - TimesTen setup, 4-4
 - using native gateways, 15-2
 - XML, 4-19
- data types, transforming, 10-6
- database
 - allowing or disallowing execution privileges, 13-13
 - features, about, 6-4
 - features, for driving tables, 8-10
 - row-level security in, 13-6
- database functions
 - about, C-65
 - EVALUATE, C-66
 - EVALUATE_AGGR, C-67
 - EVALUATE_ANALYTIC, C-67
 - EVALUATE_PREDICATE, C-68
- database hints
 - about, 7-41
 - about entering SQL comment markers, 7-43
 - creating, 7-42
 - index hint, about, 7-42
 - Leading hint, about, 7-42
 - performance considerations, 7-42
 - Physical layer objects that accept hints, 7-41
 - usage examples, 7-42
- database objects
 - about database types, 6-1
 - assigning ODBC type, 6-2
 - automatically assigning type, 6-1
 - creating for standby database configuration, 4-35
 - creating manually, 6-2
 - setting up, 6-1
 - specifying features for, 6-4
 - Virtual Private Database option, 6-2
- DATABASE system function, about, C-76
- databases, supported, 1-16
- DataDirect Connect ODBC drivers, using, 15-6

DATE data type, changing to, C-61
 DATE_DISPLAY_FORMAT configuration option, 2-17
 DATE_TIME_DISPLAY_FORMAT configuration option, 2-17
 datetime literals, about, C-18
 DAY_OF_QUARTER calendar date/time function, about, C-55
 DAYNAME calendar date/time function, about, C-55
 DAYOFMONTH calendar date/time function, about, C-55
 DAYOFWEEK calendar date/time function, about, C-55
 DAYOFYEAR calendar date/time function, about, C-56
 DB2 CLI (Unicode) connection protocol, 4-12
 DB2 Connect, on IBM z/OS and s/390 platforms, 15-13
 DB2 Cube Views
 deploying cube metadata, F-15
 guidelines for materialized query tables (MQTs), F-16
 DB2 Cube Views Generator
 about, F-1, F-14
 about input file for, F-4
 about output files for, F-5
 conversion rules, F-6
 exporting metadata into DB2 Cube Views, F-14
 generating import file, F-1
 optional parameters and defaults, F-3
 troubleshooting, F-5
 DBFeatures.INI file, 6-4
 decimal literal, about, C-18
 Default logging level option, 2-5
 default prefix, for aggregate persistence, 12-2
 DEFAULT_PRIVILEGES configuration option, 1-15, 2-17
 deferring execution of initialization blocks, 18-10
 DEGREES math function, about, C-49
 delegated administration in multiuser development, A-18
 deleting objects, 2-13
 Density option, for Oracle OLAP, 7-33
 derived columns, creating, 8-13
 Derived from existing columns using an expression option, 8-13
 DESCRIPTION system session variable, 18-5
 descriptor ID column, assigning to logical column, 8-13
 design guidelines for a repository, 1-10
 Diagrams
 about using, 2-10
 dimensions
 about, 1-7
 about hierarchies in, 9-1
 automatically creating, 9-10
 automatically creating level counts, 9-12
 identifying, 1-7
 with multiple hierarchies, 1-8
 dimension-specific aggregation rules, setting up, 8-16
 direct database requests, allowing or disallowing, 13-13
 DISABLE_CACHE_HIT system session variable, 18-5
 DISABLE_CACHE_SEED system session variable, 18-5
 DISABLE_PLAN_CACHE_HIT system session variable, 18-6
 DISABLE_PLAN_CACHE_SEED system session variable, 18-6
 DISABLE_SUBREQUEST_CACHE system session variable, 18-6
 Disallow Online RPD Updates configuration option, 2-17
 Disallow RPD Updates configuration option, 2-21
 Discard Local Changes menu option, 3-13
 Display Column, for Essbase cubes, 7-24
 display functions
 example, C-12
 reset behavior, C-12
 Display original names for alias in diagrams option, 2-3
 Display qualified names in diagrams option, 2-3
 DISPLAYNAME system session variable, 18-5
 documentation of repository mappings, generating, 17-13
 double column support, enabling, 8-13
 dragging and dropping
 business models, 11-2
 physical tables and columns, 8-1
 driving tables
 about, 8-10
 tuning performance, 8-10
 DSNs, setting up for the Oracle BI Server, 1-16
 duplicating business model and subject area, 8-3
 dynamic name
 for physical catalog or physical schema, 7-5
 for physical tables, 7-11
 dynamic repository variables
 about, 18-2
 initializing, 18-8

E

Edit menu, about, 2-8
 editing objects, 2-13
 Entity Objects, 5-2
 equalizerpds utility, 16-5
 equalizing objects, 16-3
 equalStringSet parameter, for equalizing objects, 16-6
 Essbase
 alternate hierarchies, 7-25
 associating member attributes, 7-25
 client libraries, 4-6, 15-12
 column types, 7-21
 defining aggregation rules for, 8-18
 flattening metadata, 7-22

- Gen 1 levels, 7-23
- Generations, 7-23
- importing from, 4-15
- importing UDAs, 4-18
- incremental import, 7-23
- measure hierarchies, 7-27
- modeling UDAs, 7-25
- setting up, 4-6, 15-12
- SSO using CSS Token, 6-13
- substitution variables, 7-22
- using unqualified member names, 7-27
- working with alias tables, 7-24
- working with in the Physical layer, 7-21
- ESSBASEPATH, setting, 4-6, 15-13
- Estimate Levels, using, 9-12
- EVALUATE database function, about, C-66
- Evaluate_Aggr aggregation rule, 8-15
- EVALUATE_AGGR database function, about, C-67
- EVALUATE_ANALYTIC database function, about, C-67
- EVALUATE_PREDICATE database function, about, C-68
- event polling
 - setting up for standby database configuration, 4-39
 - using the Oracle BI Event Tables utility, 17-10
- exchanging metadata
 - about, F-1
 - generating import file, F-1
 - with IBM DB2 using DB2 Cube Views, F-14
 - with Oracle Database using SQL Access Advisor, F-10
- execution precedence, establishing for initialization blocks, 18-16
- EXP math function, about, C-49
- Expand All option, for Diagrams, 2-11
- Export logical keys option, about using with parameterized SQL queries, 11-5
- Expression Builder
 - about using, 17-1
 - accessing, 17-1
 - building an expression, 17-6
 - categories, about, 17-4
 - example expression, 17-5
 - navigating, 17-6
 - toolbar, 17-3
- expression literals
 - character literals, about and expressing, C-17
 - decimal, about and expressing, C-18
 - integers, about and expressing, C-18
- external aggregation, 7-16
- External Expression, for ADF Business Components, 5-2
- Externalize Strings utility, using, 17-10
- EXTRACTBIT math function, about, C-50
- extractprojects utility, 3-11

F

feature table

- changing entries using Query DBMS, 6-5
- restoring default entries for, 6-5
- viewing in Features tab of Database dialog, 6-4
- file compression, 3-2
- File menu, about, 2-6
- filters, setting up, 13-3
- FIRST aggregate function, about, C-27
- Fit option, for Diagrams, 2-11
- floating point literal, about, C-19
- FLOOR math function, about, C-50
- foreign keys
 - relationship with primary keys, 7-34
- fragmentation content
 - about, 10-12
 - specifying, 10-8
 - specifying multicolumn content descriptions, 10-13
 - specifying parallel content descriptions, 10-13
 - specifying single column range-based predicates, 10-12
 - specifying single column value-based predicates, 10-12
 - specifying unbalanced parallel content descriptions, 10-15
- fragmented data, about, 7-35
- FROM clause syntax, about, C-4
- full repository merges, about, 16-8
- fully denormalized schemas, about, 1-9
- fully qualified names, displaying, 2-3
- Fusion Middleware Control
 - Disallow Online RPD Updates, 2-17
 - Repository File, 2-17
 - settings for repository builders, 2-17

G

- Gen 1 levels
 - about, 7-23
 - skipping, 2-4
- Generate ADF Label option, 5-13
- Generate ADF Tooltip option, 5-13
- Generate Metadata Dictionary utility, using, 17-14
- Generate target DDL in a separate file option, 12-3
- governors, setting up, 13-12
- grains, for time queries, 9-22
- grand total dimension hierarchy, example of, 9-8
- grand total levels, about, 9-2
- GROUP BY clause
 - query behavior with and without, C-12
 - syntax, about, C-5
- GROUP system session variable, 18-5
- GROUPBYCOLUMN aggregate function, about, C-27
- GROUPBYLEVEL aggregate function, about, C-28
- groups, managing, 1-16, 13-2

H

- hardware requirements, 1-16
- Help menu, about, 2-10

help, accessing, 1-11
Hide level based measure option, 2-5
Hide unusable logical table sources in Replace wizard option, 2-5, 17-9
hierarchies
 about, 1-7, 9-1
 about level-based, 9-2
 about multiple, 1-8
 determining type, 1-7
 example of grand total hierarchy, 9-8
 grand total levels, 9-2
 level attributes, about, 9-2
 level keys, about, 9-2
 level-based measure calculations, about, 9-7
 levels and distances in parent-child hierarchies, 9-14
 levels, creating, 9-2
 parent-child, about, 9-13
 setting up level-based measure calculations, 9-8
 skip-level, 9-3
 time dimensions, about, 9-3
 unbalanced, 9-3
hierarchy navigation functions
 about, C-69
 ISANCESTOR, C-69
 ISCHILD, C-70
 ISDESCENDANT, C-71
 ISLEAF, C-72
 ISPARENT, C-73
 ISROOT, C-74
hints
 about, 7-41
 creating, 7-42
 support for, 7-41
history for multiuser development, 3-20
HOUR calendar date/time function, about, C-56
HTML tables
 XML Gateway, accessing by, 4-28
Hyperion Financial Management
 Application Builder, 4-7
 importing from, 4-17
 POV value, 7-29
 query support, 7-29
 setting up, 4-7
 SSO using CSS Token, 6-13
 working with, 7-28

IBM DB2 Connect, configuring queries for UNIX, 15-13
IBM DB2 Cube Views, 2-16
 about, F-1
 using to create DB2 materialized query tables, F-14
icons, changing, 2-14
IDENTIFIER_QUOTE_CHAR parameter, 15-7
Identity Manager, 2-9
Identity menu option, 2-9
IFNULL conversion function, about, C-62

implicit fact column, setting, 11-4
Import Metadata Wizard, 4-11, 4-14, 4-21, 5-9
importing
 ADF Business Components, 5-2, 5-9
 data source preconfiguration tasks, 4-2
 from Essbase, 4-15
 from Hyperion Financial Management, 4-17
 from multidimensional data sources, 4-14
 from Oracle OLAP, 4-16
 from Oracle RPAS, 4-15
 from relational data sources, 4-10
 from SAP/BW, 4-17
 from XML data sources, 4-19
 from XMLA, 4-16
 locally, 4-10
 metadata, 4-1
 through the Oracle BI Server, 4-11
 XML data using ODBC, 4-30
In SQL logical operator, about, C-14
incremental import, for Essbase, 7-23
INDEXCOL conversion function, about, 17-7, C-62
indexing, about index hint instructions, 7-42
Informix Database, connecting using DataDirect Connect driver, 15-9
initialization blocks
 about connection pools for, 6-6
 about using with variables, 18-7
 allowing deferred execution of, 18-10
 associating with variables, 18-15
 creating, 18-10
 deferring execution, 18-8
 enabling and disabling, 18-17
 execution order, setting, 18-16
 initializing dynamic repository variables, 18-8
 initializing session variables, 18-8
 Required for authentication option, 18-10
 row-wise initialization, 18-8
 rules for deferring execution, 18-17
 working with, 18-7
INSERT string function, about, C-41
installation types, supported, 1-16
integers literals, about, C-18
INTERRUPT_ENABLED parameter, 15-13
Is Null SQL logical operator, about, C-14
ISANCESTOR hierarchy navigation function, about, C-69
ISCHILD hierarchy navigation function, about, C-70
ISDESCENDANT hierarchy navigation function, about, C-71
ISLEAF hierarchy navigation function, about, C-72
ISPARENT hierarchy navigation function, about, C-73
ISROOT hierarchy navigation function, about, C-74

J

JDBC data source, for Oracle ADF application, 5-8
JDKs, supported, 1-16
Job Manager, 2-8
Jobs menu option, 2-8

Joins Manager
 about, 2-9
 creating logical joins, 8-8
 creating physical joins, 7-37
 using to create joins, 8-7
Joins menu option, 2-9

K

keyboard shortcuts in Administration Tool, G-1
KPIs, 1-13

L

labels, for ADF data sources, 5-12
LAST aggregate function, about, C-28
Leading hint, about, 7-42
Leaf column type, 7-21
LEFT string function, about, C-41
LENGTH string function, about, C-42
level attributes, about, 9-2
level keys, about, 9-2
level-based hierarchies, about, 9-2
level-based measures
 about, 9-7
 calculations, 9-7
 query results, 9-7
lexicographical sorting for logical columns, 8-12
lifecycle management, for repository, A-1
Like SQL logical operator, about, C-14
limits
 limiting maximum run time, 13-13
 limiting rows received, 13-12
 restricting queries to time periods, 13-13
Linux
 about data source setup, 15-1
 connecting to data sources using native ODBC drivers, 15-10
 database setup script, locating, 15-2
 DataDirect Connect ODBC drivers, 15-6
literals, in SQL, C-17
Load all objects on startup option, about
 selecting, 2-19
Load Decision File option, 16-14
LOCALE configuration option, 2-17
localization, about, 1-16
localizing Presentation layer object names and descriptions, 17-10
LOCATE string function, about, C-42
LOCATEN string function, about, C-43
LOG math function, about, C-50
LOG10 math function, about, C-50
logging level, setting for BISystem user, 2-5
logging, managing, 1-16
logical columns
 assigning descriptor ID, 8-13
 associating with logical levels, 9-6
 changing sort order for, 8-12
 configuring for multicurrency support, 8-14
 creating, 8-11

 creating derived columns, 8-13
 enabling write back on, 8-21
 lexicographical sorting, 8-12
 mapping to physical columns, 10-5
 moving or copying, 8-20
 unmapping from source, 10-6
logical dimension tables
 about, 1-6
 identifying, 1-6
logical dimensions
 about, 9-1
 about including key column, 9-6
 associating logical columns and tables with, 9-6
 automatically populating level counts, 9-12
 creating, 9-5
 creating automatically, 9-10
 creating logical levels, 9-5
 creating time dimensions, 9-25
 creating with parent-child hierarchies, 9-17
 for multidimensional data sources, 9-1
 modeling time dimensions, 9-21
 with level-based hierarchies, 9-2
 with parent-child hierarchies, 9-13
logical display folders, setting up, 8-22
logical fact tables
 about, 1-5
 identifying, 1-5
 in a single business model, 1-13
 joins, 1-6
logical foreign key joins, 8-7, 8-8
logical joins
 about, 8-6
 about creating, 8-7
 about driving tables, 8-10
 cardinality, 8-8
 creating, 8-8
 creating with Joins Manager, 8-8
 displaying physical tables for, 8-11
 displaying with Business Model Diagram, 8-7
 tuning performance for driving table, 8-10
logical keys, exporting in Presentation layer, 11-4
logical levels
 about keys, 9-2
 associating logical columns for, 9-6
 automatically populating level counts, 9-12
 creating, 9-5
 designating chronological keys, 9-26
 grand total levels, example, 9-9
 hierarchy, about, 9-2
 populating level counts, 9-12
 primary key, adding, 9-9
 specifying preferred drill path, 9-10
logical objects, displaying mapped physical objects for, 8-11
logical operators, C-14
Logical SQL, about, C-1
logical stars and snowflakes, creating corresponding subject areas for, 11-2
logical table sources
 about setting up, 10-1

- creating, 10-1
- defining aggregate table content definitions, 10-7
- defining content of, 10-6
- limiting number of rows returned, 10-9
- managing, 10-1
- mapping columns in, 10-5
- merge algorithms for, D-4
- setting priority group numbers, 10-2
- settings for parent-child hierarchies, 10-10
- using the Replace Column or Table Wizard, 17-7
- Where clause filter, using to constrain physical tables, 10-10
- logical tables
 - adding new logical table source, 10-1
 - associating with logical dimension, 9-6
 - creating and managing, 8-4
 - creating by dragging and dropping, 8-5
 - creating manually, 8-5
 - designating as lookup tables, 8-5
 - key, specifying, 8-6
 - working with, 8-4
- LOGLEVEL system session variable, 18-5
- logs for master repository, 3-17
- lookup tables, 1-8, 2-22, 8-5
- LOWER string function, about, C-43

M

- main window, Administration Tool, 2-2
- Manage menu, about, 2-8
- ManageRepositories permission, 2-19
- managing authentication, 13-2
- mappings, creating between Physical layer and Business Model and Mapping layer, 10-1
- Marketing menu option, 2-9
- Marquee Zoom option, for Diagrams, 2-11
- master repository
 - about, 3-2
 - log, 3-17
 - tracking changes in, 3-17
- Materialization option, for Oracle OLAP, 7-33
- materialized query tables (MQTs), about, F-14
- materialized views
 - creating, F-13
 - using, F-10
- math functions
 - about, C-47
 - ABS, C-48
 - ACOS, C-48
 - ASIN, C-48
 - ATAN, C-48
 - ATAN2, C-48
 - CEILING, C-49
 - COS, C-49
 - COT, C-49
 - DEGREES, C-49
 - EXTRACTBIT, C-50
 - FLOOR, C-50
 - LOG, C-50
 - LOG10, C-50

- MOD, C-51
- PI, C-51
- POWER, C-51
- RADIANS, C-51
- RAND, C-51
- RANDFROMSEED, C-52
- ROUND, C-52
- SIGN, C-52
- SIN, C-52
- SQRT, C-53
- TAN, C-53
- TRUNCATE, C-53
- mathematical operators, about, C-15
- MAVG running aggregate function, about, C-32
- MAX aggregate function, about, C-28
- MAX_PARAMETERS_PER_DRIVE_JOIN
 - parameter, 8-10
- MAX_QUERIES_PER_DRIVE_JOIN parameter, 8-11
- maximum number of connections, 4-5, 6-9
- MaxThreadsPerClient setting, for Microsoft Analysis Services, 6-9
- measures
 - about calculations for level-based measures, 9-7
 - associating with levels in a dimension, 8-19
 - Based on dimensions option, 8-16
 - behavior with aggregate functions, C-11
 - creating time series measures, 9-27
 - externally aggregated, 7-16
 - specifying default aggregation rule for, 8-15
- MEDIAN aggregate function, about, C-29
- Member Alias column type, 7-21
- member counts, viewing, 7-20
- Member Key column type, 7-21
- memnor columns, 7-21
- memory requirements, 1-16
- menus, about, 2-6
- Merge Local Changes menu option, 3-13
- Merge Repository Wizard
 - rules for, D-1
 - using, 16-10, 16-15, 16-19, D-3
- merging repositories
 - about, 16-8
 - about full merge, 16-8
 - about patch merge, 16-16
 - best practices, A-21
 - choosing merge method, A-18
 - equalizing objects, 16-3
 - merge rules, D-1
 - with a common parent, 16-10
 - without a common parent, 16-15
- metadata
 - querying and managing, 16-21
- metadata dictionary, generating, 17-14
- Microsoft Analysis Services
 - avoiding connection errors for, 6-9
 - importing from, 4-16
 - MaxThreadsPerClient setting, 6-9
- Microsoft SQL Server, connecting using DataDirect Connect driver, 15-6
- migrating the repository, A-5

- MIN aggregate function, about, C-29
- minimum disk space requirements, 1-16
- MINUTE calendar date/time function, about, C-56
- Miscellaneous tab, in Connection Pool dialog, 6-19
- MOD math function, about, C-51
- modeling
 - alternate hierarchies, 7-25
 - best practices for presentation tables, 11-9
 - bridge tables, 8-23
 - Essbase data in the Physical layer, 7-21
 - Hyperion Financial Management data, 7-28
 - Oracle OLAP data, 7-30
 - parent-child relationship tables, 9-20
 - time series data, 9-21
 - UDAs, 7-25
- modes, offline and online, 2-18
- modified repository, about, 16-8
- MONTH calendar date/time function, about, C-56
- MONTH_OF_QUARTER calendar date/time function, about, C-56
- MONTHNAME calendar date/time function, about, C-57
- MSUM running aggregate function, about, C-32
- MUD
 - history, 3-20
 - merge, 3-14
 - shared network directory, 3-7
- multicurrency support, configuring, 8-14
- multi-database joins, about, 7-35
- multidimensional data sources
 - about, 1-9
 - importing, 4-14
- multilingual schemas, 1-8
- multiple hierarchies
 - about, 1-8
 - in the Presentation layer, 11-10
- multiuser development
 - about, 3-1
 - architecture, A-4
 - best practices, A-19
 - branching, 3-17, A-16
 - case study, B-1
 - concepts, A-4
 - delegated administration, A-18
 - governance best practices, A-2
 - history, about viewing and deleting, 3-20
 - local changes, merging, 3-13
 - making changes, 3-8
 - merge, about, 3-14
 - metadata, changing and testing, 3-12
 - options, 3-21
 - planning for deployment, A-1
 - platforms, A-10
 - projects, 3-3, 3-9, 3-13, A-14
 - rolling back, A-26
 - sandboxes, A-4
 - setting directory, 2-6
 - setting up environment, about, 3-7
 - shared network directory, 3-7
 - styles, A-6

- troubleshooting, A-24
- understanding, A-13
- Multiuser development directory option, 3-8

N

- naming restrictions for repository objects, 2-13
- NATURAL_JOIN keyword, using in SELECT_ PHYSICAL statements, C-8
- nesting presentation tables, 11-7
- net service name, using for Oracle Database connections, xxxi
- New Join option, for Diagrams, 2-12
- New Table option, for Diagrams, 2-12
- nonsystem session variables, 18-6
- normalized schemas, about, 1-9
- NOW calendar date/time function, about, C-57
- nqcmd utility
 - disallowing access, 6-3
 - running for cache seeding, 6-3
 - using, 14-3
- NQSConfig.INI
 - AGGREGATE_PREFIX, 12-2
 - DATE_DISPLAY_FORMAT, 2-17
 - DATE_TIME_DISPLAY_FORMAT, 2-17
 - DEFAULT_PRIVILEGES, 1-15, 2-17
 - LOCALE, 2-17
 - NULL_VALUES_SORT_FIRST, 4-19
 - options for repository builders, 2-17
 - PREVENT_DIVIDE_BY_ZERO, 8-13
 - TIME_DISPLAY_FORMAT, 2-17
 - using, 1-16
- NTILE aggregate function, about, C-29
- NULL_VALUES_SORT_FIRST configuration option, 4-19
- nullable, physical columns, 7-13
- Number of Cache Entries option (authentication cache), 2-6
- Number of elements at this level option, 9-6
- numeric literals, about, C-18

O

- OBIEEBroker
 - deploying as a shared library, 5-4
 - URL for, 5-10, 6-20
- object permissions, setting up, 13-8
- objects
 - checking in, 2-20
 - checking out, 2-20
 - deleting, 2-13
 - editing, 2-13
 - naming requirements for, 2-13
 - reordering, 2-13
 - selecting, 2-13
 - sorting, 2-15
- OCI 10g/11g connection protocol, 4-12
- OCI connections, troubleshooting, 15-4
- OCTET_LENGTH string function, about, C-43
- ODBC

- 2.0 and 3.5 connection protocol, 4-12
 - about query failures, 4-5
 - data sources, importing metadata from, 4-1
 - DataDirect Connect drivers on Linux and UNIX, 15-6
 - native drivers, 15-10
 - ODBC DSN
 - for the Oracle BI Server, 1-16
 - setting up for data sources, 4-3
 - odbc.ini file, 15-7, 15-8, 15-9, 15-10, 15-11
 - offline mode, 2-18
 - online help, accessing, 1-11
 - online mode, 2-19
 - opaque views
 - about, 7-38
 - deleting, guidelines for, 7-41
 - deploying, 7-38
 - redeploying, guidelines for, 7-41
 - undeploying, 7-40
 - when to use, 1-12
 - operating systems, supported, 1-16
 - operators, in SQL, C-14
 - Options dialog, using, 2-3
 - Oracle Application Development Framework,
 - about, 5-1
 - Oracle BI Event Tables utility, using, 17-10
 - Oracle BI Scheduler setup, for standby database configuration, 4-39
 - Oracle BI Server
 - architecture, 1-1
 - nonlocal files, about accessing, 4-20
 - ODBC DSN configuration, 1-16
 - starting and stopping, 1-15
 - utilities, 2-15
 - Oracle BI Server Web services, using, 1-15
 - Oracle BI Server XML API
 - using, 1-15, 2-16
 - using to update connection pool settings, 6-7
 - Oracle BI Server XML Gateway
 - See XML Gateway
 - Oracle Business Intelligence
 - home page on OTN, 2-10
 - localizing, 1-16
 - logging, 1-16
 - new features, xxiii
 - starting and stopping, 1-15
 - Oracle Database
 - about exchanging metadata with, F-1
 - deploying metadata in, F-11
 - importing from, 4-12
 - setting up, 4-4
 - Table or view not found error, 6-10
 - using net service name, xxxi
 - using stored procedures with, 7-7
 - Oracle Database Metadata Generator
 - about, F-1
 - about input file for, F-4
 - about output files for, F-5
 - conversion rules, F-6
 - exporting metadata into SQL Access Advisor, F-10
 - generating import file, F-1
 - optional parameters and defaults, F-3
 - troubleshooting, F-5
 - Oracle Database SQL Access Advisor, 2-16
 - Oracle instance, initializing shell window for, 2-16
 - Oracle Marketing Segmentation, 2-9, 2-15
 - Oracle OLAP
 - about, 7-30
 - about cubes and columns, 7-33
 - about dimensions, hierarchies, and levels, 7-31
 - default aggregation rule, 7-31
 - Density option, 7-33
 - importing from, 4-16
 - Materialization option, 7-33
 - setting up, 4-4
 - working with, 7-30
 - Oracle RPAS
 - about, 4-18
 - importing from, 4-15
 - recommended metadata types for import, 4-18
 - setting SQLExtendedFetch option, 4-10
 - setting up, 4-10, 15-11
 - Oracle Scorecard and Strategy Management
 - enabling comments and status overrides for, 14-1
 - modeling requirements for, 1-13
 - Oracle WebLogic Server, managing, 1-16
 - OracleADF_HTTP connection type, 5-9
 - ORDER BY clause syntax, C-5
 - original repository, about, 16-8
 - orphan locks, about, A-24
 - Other column type, 7-21
 - outer joins, modeling, 1-14
 - Outline Sort column type, 7-21
- ## P
-
- paint.rpd, 2-18
 - Pan option, for Diagrams, 2-11
 - parallel content descriptions, examples, 10-14
 - Parent Reference column type, 7-21
 - parent-child hierarchies
 - about, 9-13
 - about levels and distances, 9-14
 - creating dimensions for, 9-17
 - logical table source settings, 10-10
 - maintaining, 9-21
 - parent-child relationship tables
 - about, 9-15
 - defining, 9-18
 - modeling, 9-20
 - patch repository merges
 - about, 16-16
 - applying patches, 16-19
 - generating patches, 16-18
 - patchrpd utility, about, 16-20
 - PERCENTILE aggregate function, about, C-29
 - performance
 - accelerating by exchanging metadata with databases, F-1

- avoiding extra physical join conditions, 7-36
- avoiding multi-database joins for, 7-35
- best physical joins for, 7-34
- considerations for driving tables, 8-10
- defining aggregations correctly, 8-14
- defining chronological key at query grain, 9-22
- improving using aggregate tables, 12-1
- improving using database hints, 7-42
- improving using query caching, 1-16
- improving using unqualified member names, 7-27
- improving with standby database configuration, 4-34
- indexing relational tables when persisting aggregates, xxix
- removing unnecessary objects from Physical layer, 4-14
- setting correct aggregation rules, 7-15
- using OCI to connect to Oracle Database, 4-12
- PERIODROLLING time series function, about, C-37
- permission inheritance in the repository, 13-10
- permissions
 - default, 13-2
 - for connection pools, 6-8
 - for Presentation layer objects, 11-12
 - limiting queries by setting up data filters, 13-4
 - report for Presentation layer objects, 11-14
 - required for online mode, 2-19
 - sorting columns, 11-14
- PERMISSIONS system session variable, 18-5
- persist connection pool
 - in Database dialog, 6-2
 - setting up, 6-21
- physical catalogs
 - creating, 7-4
 - specifying dynamic name, 7-5
- physical columns
 - creating or editing, 7-12
 - nullable, 7-13
 - row counts for, 7-43
 - type, 7-13
 - viewing data, 7-14
- physical connection settings, updating, 3-12
- physical cube tables, about, 7-15
- Physical Diagram
 - about using, 2-10
 - about using to specify multi-database joins, 7-35
 - accessing from Business Model and Mapping layer, 8-11
 - displaying, 7-2
 - displaying qualified names in, 2-3
 - foreign key join or complex join, defining, 7-36
 - physical joins, about defining, 7-36
 - working with, 7-2
- physical dimensions, about, 7-17
- physical display folders, setting up, 7-5
- physical hierarchies
 - about, 7-17
 - adding or removing cube columns, 7-18
 - setting the hierarchy type, 7-18
- physical joins
 - about, 7-34
 - about complex joins, 7-34
 - about imported key and foreign key joins, 7-34
 - about primary key and foreign key relationships, 7-34
 - avoiding unnecessary joins, 7-36
 - creating with Joins Manager, 7-37
 - defining in the Physical Diagram, 7-36
 - fragmented data, about, 7-35
 - multi-database joins, about, 7-35
 - working with, 7-1
- Physical layer
 - about, 1-3
 - creating and maintaining, 4-1
 - creating manually, 6-2
 - hints, about, 7-41
 - identifying source content, 1-8
 - mapping logical columns to physical columns, 10-4
 - modeling bridge tables, 8-24
 - physical joins, defining with the Joins Manager, 7-37
 - queries, specifying types sent to a database, 6-4
 - removing unused objects, 17-15
 - setting up display folders, 7-5
 - tips for, 1-11
 - updating objects in, 17-12
 - working with dimensions and hierarchies, 7-15
 - working with Essbase data sources, 7-21
 - working with Hyperion Financial Management data, 7-28
 - working with multidimensional sources, 7-15
 - working with Oracle OLAP data, 7-30
 - working with physical tables, 7-6
- physical schemas
 - about importing, 4-1
 - creating, 7-5
 - specifying dynamic name, 7-5
- physical tables
 - about, 7-6
 - cache settings, 7-11
 - creating or editing, 7-10
 - row counts for, 7-43
 - setting XML properties for, 7-12
 - specifying columns and keys for, 7-12
 - table types, 7-6
 - viewing data, 7-14
 - virtual physical tables, creating, 7-6
 - working with, 7-1, 7-6
- PI math function, about, C-51
- platforms, supported, 1-16
- Point of View (POV) value, for Hyperion Financial Management data sources, 7-29
- policies, managing, 13-2
- Populate privilege, allowing or disallowing, 13-14
- POPULATE SQL, 6-3
- PORTALPATH system session variable, 18-5
- POSITION string function, about, C-44
- POWER math function, about, C-51

- preferences, setting, 2-3
- preferred drill path, identifying, 9-10
- PREFERRED_CURRENCY session variable, 8-14
- presentation columns
 - about working with, 11-7
 - Alias tab, using, 11-14
 - creating, 11-7
 - removing, 11-3
 - renaming, 11-4
 - reordering, 11-8
- presentation hierarchies
 - about, 11-8
 - based on logical dimensions with multiple hierarchies, 11-10
 - creating and managing, 11-9
 - editing properties for, 11-11
- Presentation layer
 - about, 1-3, 11-1
 - Alias tab, using, 11-14
 - creating, 11-1
 - localizing names, 17-10
 - logical keys, about exporting in the subject area, 11-4
 - maintaining, 11-4
 - nested folders in Oracle BI Answers, 1-15, 11-7
 - permission report, 11-14
 - presentation columns, working with, 11-7
 - removing unnecessary columns, 11-3
 - renaming presentation columns, 11-4
 - setting permissions in, 11-12
 - tips for, 1-14
- presentation levels
 - about, 11-8
 - creating and managing, 11-11
 - specifying columns for display, 11-12
- presentation tables
 - Alias tab, using, 11-14
 - best practices for modeling, 11-9
 - creating and managing, 11-6
 - nesting, 11-7
 - reordering columns in, 11-8
 - reordering in subject area, 11-7
- PREVENT_DIVIDE_BY_ZERO configuration option, 8-13
- primary key
 - foreign key, relationship with, 7-34
 - specifying, 7-14
- priority group numbers
 - about, 10-2
 - example, 10-3
- processes, starting and stopping, 1-15
- production environments, moving to, 1-16
- Project Manager, 2-9
- Project menu option, 2-9
- projects
 - about, 3-4, A-14
 - best practices, A-20
 - checking in, 3-15
 - checking out, 3-9
 - creating, 3-5

- tracking checkouts and checkins, 3-17
- upgrading, 3-7
- Prompt when moving logical columns option, 2-4, 8-20
- PROXY system session variable, 18-4
- prunerpd utility, using, E-1
- pruning repository objects, E-1
- Publish to Network menu option, 3-13

Q

- qualified names, displaying, 2-3
- QUARTER_OF_YEAR calendar date/time function, about, C-57
- queries
 - aggregate functions, rules for, C-12
 - database, specifying types sent to, 6-4
 - limiting by maximum run time, 13-13
 - limiting by number of rows, 13-12
 - restricting to time periods, 13-13
- query caching
 - cache expiration time, 7-11
 - cache persistence time, 7-11
 - including tables for, 7-11
 - managing, 1-16
- query candidates, identifying for aggregation, 12-2
- Query DBMS button, using to change Feature table entries, 6-5
- query limits, setting, 13-12
- Query Related Objects feature, 16-25
- Query Repository menu option, 2-10
- Query Repository tool, 16-21
- query workload, creating, F-12
- querying repository metadata
 - about, 16-21
 - filtering results, 16-23
 - using Query Related Objects, 16-25
 - using Query Repository tool, 16-21

R

- RADIANS math function, about, C-51
- ragged hierarchies
 - See unbalanced hierarchies
- RAND math function, about, C-51
- RANDFROMSEED math function, about, C-52
- RANK aggregate function, about, C-30
- RCOUNT running aggregate function, about, C-34
- rdaadmin client tool, 15-12
- read-only mode, about, 2-21
- refresh repository view, 2-8
- relational data sources
 - importing, 4-10
 - schema types, 1-9
- Remove Unused Physical Objects utility, using, 17-15
- Remove unused physical tables after Merge option, 2-4
- rename map, for equalizing objects, 16-6
- Rename Wizard, using, 17-11

- reordering objects, 2-13
- REPEAT string function, about, C-44
- Replace Column or Table Wizard, using, 17-7
- REPLACE string function, about, C-44
- repository
 - applying patches, 16-19
 - architecture, 1-3
 - checking consistency, 14-2
 - comparing with another repository, 16-1
 - compressed format, 3-2
 - design guidelines, 1-10
 - documentation for mappings, 17-13
 - equalizing objects, 16-3
 - generating patches, 16-18
 - layers, 1-3
 - making available for queries, 14-5
 - managing lifecycle for, A-1
 - merging with another repository, 16-8
 - migration, A-5
 - object naming restrictions, 2-13
 - object permissions, 13-8
 - offline mode, 2-18
 - online mode, 2-19
 - opening, 2-7
 - original, modified, and current, 16-8
 - pruning unwanted objects, E-1
 - query limits, 13-12
 - querying, 16-21
 - read-only mode, 2-21
 - refreshing, 2-8
 - rolling back, A-26
 - row-level security, 13-3
 - saving, 14-2
 - testing, 14-3
 - upgrading, xxix
 - uploading, 14-5
 - using variables in, 18-1
- repository builders
 - prerequisite knowledge, 1-1
 - tips for, 1-10
- Repository Documentation utility, using, 17-13
- repository file
 - blank, 16-15
 - create new, 4-2
 - options when saving, 2-21
- Repository File configuration option, 2-17
- repository objects
 - equalizing, 16-3
 - permissions, setting, 11-12
 - renaming, 17-11
- repository password, changing, 16-26
- repository variables
 - about, 18-1
 - associating with initialization blocks, 18-15
 - cache purging considerations, 18-2
 - creating, 18-3
 - dynamic, 18-2
 - initializing dynamic repository variables, 18-8
 - static, 18-2
 - using in Expression Builder, 18-3

- REQUESTKEY system session variable, 18-5
- require fully qualified table names, 6-10
- Retail Predictive Application Server
 - See Oracle RPAS
- REVERSE_LTS_PRIORITY session variable, 10-3
- REVERSED_LTS_PRIORITY_SA_VEC session variable, 10-3
- REVERSIBLE_LTS_PRIORITY_SA_VEC session variable, 10-3
- RIGHT string function, about, C-45
- RMAX running aggregate function, about, C-34
- RMIN running aggregate function, about, C-35
- ROLEGUIDS system session variable, 18-5
- ROLES system session variable, 18-5
- roles, for physical tables, 7-8
- rolling back repository to previous versions, A-26
- Root column type, 7-21
- ROUND math function, about, C-52
- row counts
 - displaying, 2-4, 7-43
 - native database, about updating in, 15-5
 - updating, 2-9, 7-44
- row-level security, setting up, 13-3
- row-wise initialization, 18-8
- RSUM running aggregate function, about, C-33
- running aggregate functions
 - about, C-31
 - MAVG, C-32
 - MSUM, C-32
 - RCOUNT, C-34
 - RMAX, C-34
 - RMIN, C-35
 - RSUM, C-33

S

- SA System subject area, about, 1-16
- sametaexport, 2-16
- sametaexport utility
 - about, F-1
 - parameters, F-2
 - running, F-2
- SampleAppLite.rpd
 - about, 2-17
 - changing password for, 16-26
 - default password for, 2-17
- Sampleapp.rpd, about, 2-17
- sandboxes, for multiuser development, A-4
- SAP/BW
 - connection pool properties, 6-14
 - importing from, 4-17
 - setting up, 4-10
- SAP/BW data sources
 - working with cube variables for, 7-19
- Save Decisions to File option, 16-14
- saving, the repository, 14-2
- schema objects, creating in Physical layer, 7-5
- schemas
 - about types of, 1-9
 - multilingual, 1-8

- physical schemas, about importing, 4-1
- relational, 1-9
- scrolling speed, 2-6
- SECOND calendar date/time function, about, C-57
- security
 - data access, 13-1
 - data filters, 13-3
 - managing, 1-16
 - object permissions, 13-8
 - permission inheritance, 13-10
 - query limits, 13-12
 - row-level, 13-3
 - summary of tasks, 13-2
- security sensitive session variables, 18-6
- Select option, for Diagrams, 2-12
- SELECT statement
 - about and basic syntax, C-2
 - conditional expressions, C-15
 - GROUP BY clause syntax, C-5
 - mathematical operators, C-15
 - ORDER BY clause syntax, C-5
 - rules for queries and aggregate functions, C-9
 - select list syntax, C-4
 - SQL logical operators, C-14
 - subquery support, C-3
 - usage notes, C-3
 - WHERE clause syntax, C-4
- SELECT_PHYSICAL statement
 - about and basic syntax, C-5
 - aggregate functions not supported in, C-6
 - queries supported by, C-7
 - usage notes, C-9
 - using the NATURAL_JOIN keyword, C-8
- SELECT_PHYSICAL system session variable, 18-6
- selecting objects, 2-13
- semantic models, independent, A-7, B-2
- Session Manager, 2-8
- session variables
 - about, 18-4
 - associating with initialization blocks, 18-15
 - creating, 18-6
 - initializing, about, 18-8
 - naming scheme for ADF data source UI hints, 5-12
 - nonsystem, 18-6
 - PREFERRED_CURRENCY, 8-14
 - REVERSE_LTS_PRIORITY, 10-3
 - REVERSED_LTS_PRIORITY_SA_VEC, 10-3
 - REVERSIBLE_LTS_PRIORITY_SA_VEC, 10-3
 - row-wise initialization, 18-8
 - security sensitive, 18-6
 - system, 18-4
 - using for authenticating users, 18-4
 - using in Dynamic Name field, 7-5
 - virtual private databases, 18-6
- Sessions menu option, 2-8
- shared login, for connection pools, 6-11
- shared network directory, for MUD, 3-7
- Show Calculation Wizard introduction page
 - option, 2-4
- Show Consistency Checker menu option, 2-9
- Show row count in physical view option, 2-4
- Show statusbar option, 2-4
- Show tables and dimensions only under display
 - folders option, 2-5
- Show toolbar option, 2-4
- SIGN math function, about, C-52
- SIN math function, about, C-52
- size, of a repository file, 3-2
- SKIN system session variable, 18-5
- Skip Gen 1 levels in Essbase drag and drop actions
 - option, 2-4
- skip-level hierarchies, about, 9-3
- snowflake schemas, about, 1-9
- software requirements, 1-16
- sort objects options, 2-6
- sorting logical columns, 8-12
- SPACE string function, about, C-45
- SQL Access Advisor
 - about, F-1
 - using to create materialized views, F-10
- SQL Bypass database
 - about specifying, 5-3
 - setting up, 5-10, 6-21
- SQL features, specifying, 6-4
- SQL functions
 - aggregate functions, about, C-24
 - calendar date/time functions, about, C-53
 - conversion functions, about, C-61
 - database functions, about, C-65
 - expressing literals, C-17
 - hierarchy navigation functions, about, C-69
 - math functions, about, C-47
 - running aggregate functions, about, C-31
 - string functions, about, C-39
 - system functions, about, C-75
 - time series functions, about, C-35
- SQL operators, C-14
- SQL syntax and semantics
 - conditional expressions, C-15
 - FROM clause syntax, about, C-4
 - GROUP BY clause syntax, about, C-5
 - including and setting variables, C-23
 - ORDER BY clause syntax, about, C-5
 - queries and aggregate functions, rules for, C-9
 - Select list syntax, C-4
 - Select statement, about and basic syntax, C-2
 - Select usage notes, C-3
 - SQL logical operators, C-14, C-15
 - WHERE clause syntax, about, C-4
- SQLExtendedFetch option, setting for Oracle RPAS, 4-10
- SQRT math function, about, C-53
- SSL, enabling, 13-3
- SSO authentication, setting up, 13-3
- SSO using CSS Token, for Essbase and Hyperion Financial Management, 6-13
- standby database configuration
 - about, 4-34
 - creating connection pools for, 4-36

- creating database object for, 4-35
- Oracle BI Scheduler configuration for, 4-39
- setting up event polling with, 4-39
- setting up usage tracking with, 4-38
- updating write-back scripts, 4-38
- using with Oracle Business Intelligence, 4-34
- write operations for primary source, 4-34
- star schemas, about, 1-9
- starting Oracle Business Intelligence processes, 1-15
- static repository variables, 18-2
- status bar, about, 2-2
- status, enabling overrides for Oracle Scorecard and Strategy Management, 14-1
- STDDEV aggregate function, about, C-30
- STDDEV_POP aggregate function, about, C-30
- stopping Oracle Business Intelligence processes, 1-15
- string functions
 - about, C-39
 - ASCII, C-39
 - BIT_LENGTH, C-40
 - CHAR, C-40
 - CHAR_LENGTH, C-40
 - CONCAT, C-40
 - EXP, C-49
 - INSERT, C-41
 - LEFT, C-41
 - LENGTH, C-42
 - LOCATE, C-42
 - LOCATEN, C-43
 - LOWER, C-43
 - OCTET_LENGTH, C-43
 - POSITION, C-44
 - REPEAT, C-44
 - REPLACE, C-44
 - RIGHT, C-45
 - SPACE, C-45
 - SUBSTRING, C-46
 - TRIMBOTH, C-46
 - TRIMLEADING, C-46
 - TRIMTRAILING, C-46
 - UPPER, C-47
- subject areas
 - Alias tab, using, 11-14
 - automatically creating based on logical stars and snowflakes, 11-2
 - creating, 11-2
 - creating multiple for a single business model, 11-2
 - duplicating with business model, 8-3
 - exporting logical keys, 11-4
 - reordering and sorting tables in, 11-7
 - setting implicit fact column for, 11-4
 - working with, 11-5
- subset repository, creating, 3-10
- substitution variables, for Essbase, 7-22
- SUBSTRING string function, about, C-46
- SUM aggregate function, about, C-31
- SUMDISTINCT aggregate function, about, C-31
- surrogate keys, 1-6

- Sybase ASE Database, connecting using DataDirect Connect driver, 15-8
- system credentials, managing, 13-2
- system functions
 - about, C-75
 - DATABASE, C-76
 - USER, C-75
- system requirements, 1-16
- system session variables, 18-4

T

- Table or view not found, 6-10
- table types, for physical tables, 7-6
- TAN math function, about, C-53
- Teradata, setting up, 4-4, 15-10
- test to production, 1-16
- testing the repository, 14-3
- text strings, using the Externalize Strings utility to translate, 17-10
- third-party products, supported, 1-16
- Tile when resizing option, 2-3
- time dimensions
 - about, 9-3
 - creating, 9-25
 - for Oracle Scorecard and Strategy Management, 1-13
 - modeling, 9-21
 - selecting and sorting chronological keys, 9-10
- time series functions
 - about, 9-21, C-35
 - AGO, 9-21, C-36
 - creating measures for, 9-27
 - grains, 9-22
 - PERIODROLLING, 9-21, C-37
 - TODATE, 9-21, C-38
 - using in Expression Builder, 9-21
- TIME_DISPLAY_FORMAT configuration option, 2-17
- TIMESTAMPADD calendar date/time function, about, C-57
- TIMESTAMPDIFF calendar date/time function, about, C-59
- TimesTen
 - avoiding query failures with, 4-5
 - setting up, 4-4
- TIMEZONE system session variable, 18-6
- tips, for repository builders, 1-10
- title bar, about, 2-2
- tnsnames.ora
 - for Oracle Database setup on UNIX, 15-2
 - for the Oracle BI Server, 4-4
 - locating with Oracle Business Intelligence environment, xxxi
- TO_DATETIME conversion function, about, C-64
- TODATE time series function, about, C-38
- toolbar
 - about, 2-2
 - options for Diagrams, 2-10
- Tools menu, about, 2-9

- tooltips, for ADF data sources, 5-12
- TOPN aggregate function, about, C-31
- transaction boundary, setting up, 6-22
- TRIMBOTH string function, about, C-46
- TRIMLEADING string function, about, C-46
- TRIMTRAILING string function, about, C-46
- troubleshooting
 - DB2 Cube Views Generator, F-5
 - Oracle Database Metadata Generator, F-5
- TRUNCATE math function, about, C-53
- Turn off Compare Mode option, 2-7, 16-3

U

- UDA column type, 7-21
- UDAs
 - importing, 4-18
 - modeling, 7-25
- UI hints, propagating for ADF data sources, 11-8
- unbalanced hierarchies, about, 9-3
- Undo Merge Local Changes menu option, 3-13
- Unicode databases, 6-19
- UNIX
 - about data source setup, 15-1
 - connecting to data sources using native ODBC drivers, 15-10
 - DataDirect Connect ODBC drivers, 15-6
 - IBM DB2 Connect for queries, configuring, 15-13
 - native databases, about updating row counts, 15-5
- unmapping logical columns, 10-6
- unqualified member names, using for Essbase, 7-27
- Update All Row Counts menu option, 2-9
- Update Physical Layer Wizard, using, 17-12
- upgrade
 - considerations, xxiv, xxix
 - IDs, 16-3
 - repositories, xxix
- UPPER string function, about, C-47
- URL
 - for ADF Business Components, 5-10, 6-20
 - for connection pools, 6-13
 - refresh interval, for XML data sources, 6-17
- usage tracking
 - managing, 1-16
 - setting up for standby database configuration, 4-38
- Use Logical Column Name property, 11-4
- user authentication, setting up, 13-2
- user interface hints, for ADF data sources, 5-11
- USER system function, about, C-75
- USER system session variable, 18-4
- USERGUID system session variable, 18-5
- USERLOCALE system session variable, 18-5
- users, managing, 1-16, 13-2
- user.sh file, 15-2, 15-6, 15-8, 15-9, 15-10
- utilities
 - Aggregate Persistence, 17-15
 - command-line, 2-15
 - equalizerpds, 16-5

- Externalize Strings, 17-10
- extractprojects, 3-11
- Generate Metadata Dictionary, 17-14
- nqcmd, 14-3
- Oracle BI Event Tables, 17-10
- patchrpd, 16-20
- prunerpd, E-1
- Remove Unused Physical Objects, 17-15
- Rename Wizard, 17-11
- Replace Column or Table in Logical Table Sources, 17-7
- Repository Documentation, 17-13
- Update Physical Layer, 17-12
- validaterpd, 2-23
- Utilities menu option, 2-10

V

- validaterpd utility, 2-23
- Value measure, for Hyperion Financial Management data sources, 7-28
- VALUEOF conversion function, about, C-65
- Variable Manager, 2-9
- variables
 - about repository variables, 18-1
 - about session variables, 18-4
 - associating with initialization blocks, 18-15
 - creating repository variables, 18-3
 - creating session variables, 18-6
 - including and setting in SQL, C-23
 - using repository variables in Expression Builder, 18-3
- Variables menu option, 2-9
- View Data, 7-14
- view links, 5-2
- View menu, about, 2-8
- View Objects, 5-2
- virtual physical tables, creating using Table Type, 7-6
- Virtual Private Database, 6-2, 13-8

W

- Web services, for Oracle BI Server, 1-15
- WEBGROUPS system session variable, 18-5
- WebLogic Domain, for ADF Business Components, 5-4
- WEEK_OF_QUARTER calendar date/time function, about, C-60
- WEEK_OF_YEAR calendar date/time function, about, C-60
- Weight Factor, for bridge tables, 8-23
- WHERE clause filters, in aggregate navigation, 10-10
- WHERE clause syntax, about, C-4
- Window menu, about, 2-10
- wizards
 - Aggregate Persistence Wizard, 12-3, 17-15
 - Calculation Wizard, 2-4, 17-15
 - Create New Repository Wizard, 4-2

- Import Metadata Wizard, 4-11, 4-14, 4-21, 5-9
- Merge Repository Wizard, 16-10, 16-15, 16-19, D-3
- Rename Wizard, 17-11
- Replace Column or Table Wizard, 17-7
- Update Physical Layer Wizard, 17-12
- write back, enabling on logical columns, 8-21
- write-back properties, for connection pool, 6-17
- write-back scripts, updating for standby database configuration, 4-38

X

- XML data sources
 - about URL for, 4-19
 - importing from, 4-19
 - setting connection pool properties for, 6-16
 - setting physical table properties for, 7-12
 - URL refresh interval, 6-17
- XML Gateway
 - about using, 4-20
 - examples, 4-23
 - HTML tables, accessing, 4-28
 - refresh interval, 4-22
 - supported security modes, 4-22
- XML ODBC
 - about using, 4-29
 - example, 4-30
 - importing XML data, 4-30
- XML utilities, 2-16
- XMLA, importing from, 4-16

Y

- YEAR calendar date/time function, about, C-61

Z

- Zoom In option, for Diagrams, 2-11
- Zoom Out option, for Diagrams, 2-11