

Oracle Database 11g: Administer a Data Warehouse

Volume I • Student Guide

D70064GC10

Edition 1.0

July 2008

D55424

ORACLE®

Authors

Lauran K. Serhal

Mark Fuller

Technical Contributors and Reviewers

Hermann Baer

Kenji Baugham

Maria Billings

Herbert Bradbury

Yanti Chang

Gerlinde Frenzen

Joel Goodman

Uwe Hesse

Yash Jain

Sean Kim

Paul Lane

Lakshmi Narapareddi

Sujatha Patchamatla

Brian Pottle

Paul Rayner

Raza Siddiqui

S Matt Taylor

Jenny Tsai

Jean-Francois Verrier

Andreas Walter

James Womack

Editors

Arijit Ghosh

Amitha Narayan

Graphic Designer

Priya Saxena

Publisher

Jobi Varghese

Copyright © 2008, Oracle. All rights reserved.

Disclaimer

This document contains proprietary information and is protected by copyright and other intellectual property laws. You may copy and print this document solely for your own use in an Oracle training course. The document may not be modified or altered in any way. Except where your use constitutes "fair use" under copyright law, you may not use, share, download, upload, copy, print, display, perform, reproduce, publish, license, post, transmit, or distribute this document in whole or in part without the express authorization of Oracle.

The information contained in this document is subject to change without notice. If you find any problems in the document, please report them in writing to: Oracle University, 500 Oracle Parkway, Redwood Shores, California 94065 USA. This document is not warranted to be error-free.

Restricted Rights Notice

If this documentation is delivered to the United States Government or anyone using the documentation on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

The U.S. Government's rights to use, modify, reproduce, release, perform, display, or disclose these training materials are restricted by the terms of the applicable Oracle license agreement and/or the applicable U.S. Government contract.

Trademark Notice

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

Contents

Preface

1 Introduction

- Objectives 1-2
- Questions About You 1-3
- Lesson Agenda 1-4
- Course Objectives 1-5
- Prerequisites and Suggested Prerequisites 1-6
- Suggested Course Schedule 1-7
- Lesson Agenda 1-8
- Sample Schemas Used in the Course 1-9
 - The Human Resources (HR) Schema 1-10
 - The Sales History (SH) Schema 1-11
- Class Account Information 1-13
- Appendixes in the Course 1-14
- SQL Environments Available in the Course 1-15
- Solutions, Labs, and Code Example Scripts 1-16
- Lesson Agenda 1-17
- Oracle 11g Useful Documentation 1-18
- Additional Resources 1-19
- Continuing Your Education: Recommended Follow-Up Classes 1-20
- Oracle Technology Network (OTN) 1-21
- Oracle by Example (OBE): <http://www.oracle.com//technology/obe> 1-22
- Technical Support: Oracle *MetaLink*: <https://metalink.oracle.com/metalink> 1-23
- Summary 1-24
- Practice 1: Overview 1-25

2 Data Warehouse Design: Overview

- Objectives 2-2
- Lesson Agenda 2-3
- Characteristics of a Data Warehouse 2-4
- OLTP Systems Versus Data Warehouses 2-5
- Data Warehouse Architectures: Basic Data Warehouse 2-7
- Data Warehouse Architectures: Basic Data Warehouse with Staging Area 2-8

Data Warehouse Architectures: Basic Data Warehouse with
Staging Area and Data Marts 2-9

Lesson Agenda 2-10

Data Warehouse Design 2-11

Data Warehouse: Design Phases 2-12

Data Warehouse Physical Design 2-13

Data Warehouse Physical Structures 2-14

Data Warehousing Objects 2-15

Characteristics of Fact Tables 2-16

Dimensions and Hierarchies 2-17

Using Hierarchies for Drill on Data and Aggregate Data 2-19

Data Warehousing Schemas 2-20

Schema Characteristics 2-21

Star Schema Model: Central Fact Table and Denormalized Dimension Tables 2-23

Star Dimensional Schema: Advantages 2-24

Snowflake Schema Model 2-25

Lesson Agenda 2-27

Optimizing Star Queries: Tuning Star Queries 2-28

Optimizing Star Queries: Star Transformation 2-29

Star Query: Example 2-30

Execution Plan Without Star Transformation 2-31

Star Transformation Prerequisites 2-32

Star Transformation Considerations 2-33

Star Transformation Rewrite Example: Phase 1 2-34

Retrieving Fact Rows from One Dimension: Phase 1 2-35

Retrieving Fact Rows from All Dimensions: Phase 1 2-36

Joining the Result Set with Dimensions: Phase 2 2-37

Star Transformation Explain Plan: Example 2-38

Star Transformation Further Optimization 2-39

Star Transformation Hints: Transaction Level 2-40

Using Bitmap Join Indexes 2-41

Star Query Optimization Versus Bitmap Joined Index Optimization 2-42

Bitmap Join Indexes: Join Model 1 2-43

Bitmap Join Indexes: Join Model 2 2-44

Bitmap Join Indexes: Join Model 3 2-45

Bitmap Join Indexes: Join Model 4 2-46

Summary 2-47

Practice 2 Overview: Using Star Transformation 2-48

3 Partitioning Concepts

Objectives	3-2
Lesson Agenda	3-3
Partitioned Tables and Indexes	3-4
Table Versus Index Partitioning	3-5
Partitioning Strategies: Single-Level Partitioning	3-6
Partitioning Strategies: Composite Partitioning	3-8
SQL Access Advisor: Overview	3-9
SQL Access Advisor: Recommendations	3-10
Partition Key	3-11
Partition Key Extensions	3-12
Lesson Agenda	3-13
Benefits of Partitioning	3-14
Performance Benefits: Partition Pruning	3-15
Manageability Benefits: Partition-Level Management	3-16
Manageability Benefits: Rolling Window Operations	3-17
Equipartitioning	3-18
Join Benefits: Full Partitionwise Joins	3-19
Join Benefits: Partial Partitionwise Joins	3-20
Lesson Agenda	3-21
Creating Partitions with Enterprise Manager	3-22
The CREATE TABLE Statement with Partitioning	3-23
Logical and Physical Attributes	3-24
Partition Strategy Declaration: Single-Level Partitioning	3-25
Partition Strategy Declaration: Composite Partitioning	3-26
Specifying Partition Attributes	3-27
The CREATE TABLE Statement with Composite Partitioning	3-28
Composite Partitioning Using a Template	3-29
Composite Partitioning and Template: Example	3-30
Partition-Extended Table Names	3-31
Partitioned Indexes	3-32
Partitioned Index Attributes: Prefixed or Nonprefixed	3-33
Global Index: Example	3-34
Local Prefix Index: Examples	3-35
Lesson Agenda	3-36
Static Partition Pruning and Star Query	3-37
Static Partition Pruning Plan: Example	3-38
Dynamic Partition Pruning and Star Query	3-39
Dynamic Partition Pruning Determination	3-40
Summary	3-41
Practice 3 Overview: Working with Partitioning Concepts	3-42

4 Parallelism Concepts

- Objectives 4-2
- Lesson Agenda 4-3
- Introduction to Parallel Execution 4-4
- Operations That Can Be Parallelized 4-5
- Scanning a Table in Parallel 4-6
- Parallel Execution with Real Application Clusters (RAC) 4-7
- Querying `GV$` Views in RAC 4-8
- Lesson Agenda 4-9
- The Granule 4-10
- Degree of Parallelism (DOP) 4-11
- Default Degree of Parallelism 4-13
- Parallel Operations 4-14
- How Parallel Execution Servers Communicate 4-15
- Lesson Agenda 4-16
- Parallel Execution Server Pool 4-17
- Minimum Number of Parallel Execution Servers 4-18
- Adaptive Multiuser and DOP 4-19
- `PARALLEL_IO_CAP_ENABLED` 4-20
- Lesson Agenda 4-21
- Enabling Parallel `DML`, `DDL`, and `QUERY` 4-22
- Enabling Parallelization and Determining DOP 4-24
- Using Parallelization Hints 4-25
- `PARALLEL` Clause: Examples 4-26
- Object's `PARALLEL` Clause 4-27
- Lesson Agenda 4-28
- Parallel Execution Plan 4-29
- `OTHER_TAG` Column 4-30
- Serial and Parallel Execution Plans 4-31
- Parallel Plan Interpretation 4-32
- Summary 4-37
- Practice 4 Overview: Understanding Parallel Concepts 4-38

5 Parallel Operations in Data Warehouses

- Objectives 5-2
- Lesson Agenda 5-3
- Parallel Query 5-4
- Parallel Partitioned Table Scan 5-5
- Parallel Partitioned Index Range Scan 5-6
- Partitionwise Joins 5-7

Lesson Agenda	5-8
Parallel DDL	5-9
Creating Indexes in Parallel	5-10
Parallel DDL: Example	5-13
Parallel DML: Overview	5-14
Performance Benefits of Parallel DML	5-15
Enabling Parallel DML	5-16
Parallel DML: Example	5-17
When to Use Parallel DML	5-18
Restrictions on Parallel DML	5-20
Lesson Agenda	5-21
Tuning Parameters for Parallel Execution	5-22
Using Default Parameter Settings	5-23
Balancing the Workload	5-24
Resource Manager and the DOP	5-26
Lesson Agenda	5-28
Execution Problems	5-29
Data Distribution and V\$PQ_TQSTAT	5-30
Modifying Data Distribution	5-32
Using V\$PX_PROCESS_SYSTAT	5-33
Using Other Dynamic Performance Views	5-34
Using V\$PX_SESSION	5-36
Using V\$PX_SESSTAT	5-37
Using V\$PX_PROCESS	5-38
Using V\$SYSSTAT	5-39
Tuning PDML	5-41
Summary	5-42
Practice 5 Overview: Performing Parallel Operations	5-43
6 Extraction, Transformation, and Loading (ETL): Extraction and Transportation	
Objectives	6-2
Lesson Agenda	6-3
What Is ETL?	6-4
Extraction Methods	6-6
Logical Extraction Methods	6-7
Physical Extraction Methods	6-8
Implementing Methods of Extraction	6-10
Capturing Change Data Without Change Data Capture (CDC)	6-11
Incremental Extraction Using CDC	6-12
Publish and Subscribe Model: The Publisher	6-13

Publish and Subscribe Model: The Subscriber 6-14
Synchronous CDC 6-16
Asynchronous CDC 6-17
Asynchronous AutoLog Mode 6-18
Asynchronous HotLog Mode 6-20
Asynchronous Distributed HotLog Mode 6-21
Factors That Influence Choice of Change Data Capture Mode 6-22
Preparing to Publish Change Data 6-24
Asynchronous Distributed HotLog Publishing: Overview 6-25
Lesson Agenda 6-28
Transportation in Data Warehouses: Overview 6-29
Transportation in a Data Warehouse 6-30
Transportable Tablespaces 6-31
Transportable Tablespaces: Example 6-32
Summary 6-34
Practice 6 Overview: Using Extract and Transport Methods 6-35

7 Extraction, Transformation, and Loading (ETL): Loading

Objectives 7-2
Lesson Agenda 7-3
Data-Loading Mechanisms 7-4
Using SQL*Loader 7-5
Using SQL*Loader: Example 7-6
Lesson Agenda 7-7
Using External Tables 7-8
Benefits of Using External Tables 7-9
Creating External Tables 7-10
Creating and Loading an External Table Using `ORACLE_LOADER`: Example 7-11
Creating and Loading an External Table Using `ORACLE_DATAPUMP`: Example 7-14
Privileges Required to Access External Tables 7-15
Defining External Tables Using SQL*Loader 7-16
Data Dictionary Information for External Tables 7-17
Changing External Data Properties 7-18
Lesson Agenda 7-19
Other Loading Methods 7-20
Summary 7-21
Practice 7 Overview: Loading Data with SQL*Loader and External Tables 7-22

8 Extraction, Transformation, and Loading (ETL): Transformation

- Objectives 8-2
- Lesson Agenda 8-3
- Data Transformation 8-4
- Multistage Data Transformation 8-5
- Pipelined Data Transformation 8-6
- Lesson Agenda 8-7
- Transformation Mechanisms 8-8
- Transformation Using SQL 8-9
- CREATE TABLE ... AS SELECT 8-10
- INSERT /*+ APPEND */ ... SELECT 8-11
- MERGE Statement: Overview 8-13
- Data Warehousing MERGE: Example 8-14
- Data Maintenance with MERGE/DELETE 8-15
- Overview of Multitable INSERT Statements 8-16
- Example of Unconditional ALL INSERT 8-17
- Example of Conditional ALL INSERT 8-18
- Example of Conditional FIRST INSERT 8-19
- Example of Mixed Conditional and Unconditional FIRST INSERT 8-20
- Example of Pivoting INSERT 8-21
- Lesson Agenda 8-22
- Overview of Table Functions 8-23
- Creating Object Types for Table Functions 8-24
- Creating Table Functions 8-25
- Using Table Functions 8-26
- Lesson Agenda 8-27
- DML Error Logging: Overview 8-28
- DML Error Logging Concepts 8-29
- Error-Logging Table 8-31
- Error-Logging Table Format 8-33
- Inserting into a Table with Error Logging: Example 8-34
- Summary 8-35
- Practice 8 Overview: Performing Transformation 8-36

9 Materialized Views

- Objectives 9-2
- Lesson Agenda 9-3
- The Need for Summary Management 9-4
- Using Summaries to Improve Performance 9-5
- Using Summaries Without Materialized Views: Example 9-6

Using Materialized Views for Summary Management 9-7
Using Materialized Views for Summary Management: Example 9-8
Summary Management Components 9-9
Overview of Materialized View (MV) Management Tasks 9-10
Planning for and Using Materialized Views 9-11
How Many Materialized Views? 9-12
One Materialized View for One Query 9-13
One Materialized View for Multiple Queries 9-14
Determining Which Materialized View to Create 9-15
Lesson Agenda 9-16
Types of Materialized Views 9-17
What Is a Nested Materialized View? 9-18
The Need for Nested Materialized Views 9-19
Using Nested Materialized Views 9-20
Nested Materialized Views: Restrictions 9-21
Cube Organized Materialized Views 9-22
Lesson Agenda 9-23
Creating a Materialized View Using the `CREATE` SQL Statement: Example 9-24
Materialized Views Storage 9-25
Estimating Materialized View Size 9-26
Specifying Build Methods 9-27
Registering Existing Materialized Views 9-28
Using Column Aliases in Materialized Views 9-29
Using Materialized View Column Alias Lists 9-30
Partitioning and Materialized Views 9-31
Partitioned Materialized View: Example 9-32
Using Enterprise Manager to Create Materialized Views 9-33
Privileges Required to Create Materialized Views 9-34
Additional Privileges Required to Create Materialized Views 9-35
Globalization and Materialized Views 9-36
Lesson Agenda 9-37
Adding Comments to Materialized Views 9-38
Altering Materialized Views 9-39
Maintaining Partitions of a Materialized View 9-41
Dropping Materialized Views 9-42
Viewing Staleness Information 9-43
Controlling the Materialized View's Integrity with the `QUERY_REWRITE_INTEGRITY`
Parameter 9-44
Invalidating Materialized Views 9-45
New and Updated MV Dictionary Views in Oracle Database 11g 9-46
Dictionary Views Showing PCT Staleness Corresponding to Base Partitions 9-47

Summary 9-49

Practice 9 Overview: Working with Materialized Views 9-50

10 Refreshing Materialized Views

Objectives 10-2

Lesson Agenda 10-3

Available Refresh Modes When Creating a Materialized View 10-4

Manual Refresh Using the `DBMS_MVIEW` Package Procedures 10-5

Using the `DBMS_MVIEW` Package: The Available `ON DEMAND`
Refresh Methods 10-6

Refresh at Scheduled Time: Using the `START WITH` and `NEXT` Clauses 10-8

Refresh Specific Materialized Views with the `REFRESH` Option 10-9

Refresh All Materialized Views with the `REFRESH_ALL_MVIEWS` Option 10-11

Refresh Dependent Materialized Views with the `REFRESH_DEPENDENT`
Option 10-12

Finding Dependent Materialized Views 10-13

When Does a Complete Refresh Occur? 10-14

Conditions That Effect Possibility of Fast Refresh 10-15

Materialized View Logs: General Concepts 10-16

Creating a Materialized View Log: Example 10-17

Restrictions for Fast-Refreshable Materialized Views 10-18

Indexes and Fast-Refreshable Materialized Views 10-21

Using the `DBMS_MVIEW` Package: Specifying the `REFRESH FORCE` Option 10-22

Partition Change Tracking (PCT) Refresh 10-24

PCT Refresh Requirements 10-25

When Is PCT Refresh Used? 10-26

Partition Key or Partition Marker? 10-27

PCT Using `ROWID` 10-29

`CONSIDER FRESH`: Fastest Refresh Scenario 10-30

`CONSIDER FRESH` Refresh Implications 10-31

Lesson Agenda 10-32

Analyzing MV Capabilities Using the `DBMS_MVIEW.EXPLAIN_MVIEW`
Procedure 10-33

Using `DBMS_MVIEW.EXPLAIN_MVIEW`: Example 10-35

General Tips for Refreshing Materialized Views 10-36

Refresh Tips for Nested Materialized Views 10-37

Summary 10-38

Practice 10: Overview 10-39

11 Working with Dimensions

- Objectives 11-2
- Lesson Agenda 11-3
- What Are Dimensions? 11-4
- What Are Dimension Objects? 11-5
- Benefits of Using Dimensions 11-6
- Lesson Agenda 11-7
- Creating Dimension Objects 11-8
- Creating Dimension Objects Using the `CREATE` Statement: Example 11-9
- Specifying the `SKIP WHEN NULL` Clause 11-10
- Dimensions Based on Multiple Tables 11-12
- Multiple Hierarchies for a Dimension: Example 11-13
- Dimensions and Privileges 11-15
- Lesson Agenda 11-16
- Viewing Dimensions Using Enterprise Manager 11-17
- Viewing Dimensions in the Data Dictionary 11-18
- Viewing the Definition of a Dimension: Using the `DESCRIBE_DIMENSION` Procedure 11-19
- Dimension Restrictions 11-20
- Verifying Relationships in a Dimension 11-21
- Lesson Agenda 11-23
- Dimension Invalidation 11-24
- Dropping Dimensions 11-25
- Constraints Versus Dimensions 11-26
- Summary 11-28
- Practice 11 Overview: Working with Dimensions 11-29

12 Query Rewrite

- Objectives 12-2
- Query Rewrite Overview 12-3
- Cost-Based Query Rewrite Process 12-4
- What Can Be Rewritten? 12-5
- Conditions Required for Oracle to Rewrite a Query 12-6
- Using Trusted Constraints 12-7
- Setting Initialization Parameters for Query Rewrite 12-8
- Query Rewrite Hints 12-10
- Will a Query Rewrite Occur? 12-11
- Using `EXPLAIN PLAN` with Query Rewrite: Example 12-12
- Using the `EXPLAIN_REWRITE` Procedure with Query Rewrite 12-13
- Using `DBMS_MVIEW.EXPLAIN_REWRITE` with an Output Table: Example 12-14

Query Rewrite Methods 12-16
Text Match Rewrite Methods 12-17
Expression Matching Query Rewrite 12-18
Full Text Matching: Example 12-19
Partition Change Tracking (PCT) and Query Rewrite 12-21
PCT and Query Rewrite: Example 12-23
Creating Indexes for Materialized Views: Recommendations 12-26
Query Rewrite Restrictions 12-27
Query Rewrite Enhancement to Support Queries Containing Inline Views 12-28
When Are Two Inline Views Equivalent? 12-29
MV IV Text Matching a Query's IV Text: Example 12-30
MV IV Text Equivalent to a Query's IV Text: Example 12-31
Transforming and Rewriting the Query from the Two Previous Examples 12-32
Query Rewrite Using Remote Tables in Oracle Database 11g 12-33
Summary 12-35
Practice 12 Overview: Using Query Rewrite 12-36

13 Using the SQL Access Advisor, Compression, and Resumable Sessions

Objectives 13-2
Lesson Agenda 13-3
What Does the SQL Access Advisor Do? 13-4
SQL Access Advisor: Overview 13-5
SQL Access Advisor: Usage Model 13-6
Possible Recommendations 13-7
SQL Access Advisor Session: Initial Options 13-8
SQL Access Advisor: Workload Source 13-10
SQL Access Advisor: Recommendation Options 13-11
SQL Access Advisor: Schedule and Review 13-13
SQL Access Advisor: Results 13-14
Lesson Agenda 13-19
SQL Access Advisor: PL/SQL Procedure Flow 13-20
SQL Access Advisor: PL/SQL Example 13-21
Privileges Needed to Use the SQL Access Advisor 13-22
Lesson Agenda 13-23
Tuning Materialized Views for Fast Refresh and Query Rewrite 13-24
Using `DBMS_ADVISOR.TUNE_MVIEW` 13-25
`TUNE_MVIEW` Output Results 13-26
Using `DBMS_ADVISOR.TUNE_MVIEW`: Example 13-28
Accessing `TUNE_MVIEW` Output Results 13-29
Lesson Agenda 13-31

Compression	13-32
Table Compression: Overview	13-33
Table Compression Concepts	13-34
What Can Be Compressed?	13-35
Using Table Compression	13-37
Compression: Examples	13-38
Maximizing Compression Ratios	13-39
Resumable Sessions	13-40
Summary	13-41
Practice 13 Overview: Using the SQL Access Advisor, Compression, and Resumable Session	13-42

Appendix A: Practices and Solutions

Appendix B: Table Descriptions

Appendix C: Using SQL Developer

Appendix D: Using SQL*Plus

Appendix E: Basic Linux and vi Commands

Appendix F: Understanding Oracle OLAP Technology

Index