



Course 272: Tuning Oracle for Data Warehousing (4 days)

Course Description...

With the emergence of data warehouses, companies are learning more about their best customers, their supply chains and product life cycles--to the benefit of the bottom line. Accurate business intelligence leads to competitive advantage, and people who know how to administer such a system play a critical role in its success

The analysis of information is playing a larger role in most organizations. In many cases, the data exists, but not in a useable form. And, when it is useable, it is not available in a meaningful timeframe.

This course extends the design and fundamental concepts of the 271 course by focusing on *how* to exploit the features of the Oracle Database Server to solve these requirements.

Learning Objectives...

- Automate aggregation with Materialized Views
- Perform statement level tuning
- Utilize bitmap and bitmap join indexes
- Partition tables and indexes
- Optimize access via the Star-Joins
- Parallel execution

Who should attend...

This course is intended for database application developers, performance tuners, system administrators, and database administrators who design, maintain, and use data warehouses.

Prerequisites...

Data Warehouse fundamentals and design, Oracle introduction and advanced SQL are necessary to fully understand this material. Courses 271, 261, and 264 plus some experience are recommended.



Course Outline

Introduction and Overview

Course Objectives

Unit 1: Statement Level Tuning

- Architecture Overview
 - Data Cache
 - Shared Pool
 - Statement Processing
- Identifying the Bottleneck
 - EXPLAIN and AUTOTRACE
 - Interpreting the output
 - SQL Trace and TKPROF
- Interpreting the results
- Optimization
 - Rules-Based vs Cost-based
 - Understanding the Execution Plan
 - Setting Optimization Goals
- Initialization and Session Settings
 - Statement level Hints
- Gathering Statistics
 - The Analyze Command
 - Using the DBMS_STATS Package
 - Identifying key table, column, and index stats

Unit 2: Indexing

- B-Trees
- B-Tree variations
- Bitmaps
- Matching the right structure with the right plan

Unit 3: Partitioning Tables and Indexes

- Defining partitions
- Range, List, and Hash Partitions
- Data Compression
- Global and Local Partitioned Indexes

Unit 4: Star Joins

- The Dimensional Design
- The Star Join Query and Optimization
- Bitmap Join Indexes



Unit 5: Parallelization

- Degree of Parallelism (DOP)
- Parallel Query
- Partition-Wise Join
- Parallel DDL and DML
- Space Management for Parallel Operations
- Parallel DML Locking Considerations
- Automated Parallel Execution Tuning Overview

Unit 6: Materialized Views

- Materialized Views Integrity
- Refresh Methods
- Partition Change Tracking (PCT) Refresh
- Refresh Modes
- Dimensions and Hierarchies
- Query Rewrite
- Summary Advisor

Unit 7: Course Summary

Summary

Please contact your ROI representative to discuss course tailoring!!!