

Oracle Database 11g: SQL Tuning Workshop (OR-DTW)

Modality: Virtual Classroom

Duration: 3 Days

About this course:

This Oracle Database 11g: (SQL Tuning Workshop R2) course helps SQL developers, DBAs and database developers in tuning and identifying ineffective SQL statements. You will explore investigative techniques to disclose different levels of insight concerning how the Oracle database performs the SQL statements; this assists you to identify the main reason for the ineffective SQL statements.

Learn To:

- How to use real-time SQL monitoring.
- Trace and control high load SQL statements.
- How to utilize Oracle tools to find ineffective SQL statements.
- Handle optimizer statistics on database objects.
- How to utilize Automatic SQL Tuning.
- Explain implementation plans, and the various ways in which data could be accessed.
- Draft the most effective SQL statements.

Benefits to You:

This course helps you to enhanced skills in relational database management as you study the method to efficiently utilize SQL commands in your business data. These functions will assist you to manipulate and query data into the database. Additionally, you will learn how to use dictionary views to recover metadata and make reports regarding their schema objects.

Study the Optimizer:

Professional trainers will also support you to learn how the optimizer selects the path. You will also study how to control the optimizer to assure the best approach is used.

Automatic SQL Tuning Tools:

This training course also incorporates the Automatic SQL Tuning resources and tools accessible in the Automatic Workload Repository. Moreover, they benefit from trace files, bind variables, and various sorts of indexes.

Note: This training course in accordance with Oracle Database 11g R 2.

Salary Estimate:

The Oracle Developer can make an average salary of \$106,225 per annum.

Course Objective:

- Describe how optimizer statistics impact the efficiency of SQL
- Track an application with its various levels of the application architecture
- Recognize inefficient performing SQL
- Know how the Query Optimizer produce decisions of how to access data
- Alter a SQL statement accomplish at its best
- Listing possible techniques of access to data, incorporating various join methods

Audience:

- Database Administrators
- Support Engineer
- Application Developers
- Data Warehouse Administrator
- Data Warehouse Developer
- SQL/PL Developer
- rehouse Administrator

Prerequisite:

- There is no prerequisite needed for this course

Course Outline:

Exploring the Oracle Database Architecture

- Oracle Database Server Architecture: Overview
- Automated SQL Execution Memory Management
- Database Storage Architecture, Logical and Physical Database Structures
- Physical Structure
- Segments, Extents, and Blocks & SYSTEM and SYSAUX Tablespaces
- Automatic Shared Memory Management
- Connecting to the Database Instance
- Oracle Database Memory Structures: Overview

Introduction to SQL Tuning

- Monitoring and Tuning Tools: Overview
- Scalability with Application Design, Implementation, and Configuration
- Simplicity in Application Design
- Reason for Inefficient SQL Performance
- Performance Monitoring Solutions

- Common Mistakes on Customer systems & Proactive Tuning Methodology
- CPU and Wait Time Tuning Dimensions
- Data Modeling, Table Design, Index Design, Using Views, SQL Execution Efficiency, Overview of SQL*Plus & SQL Developer

Introduction to the Optimizer

- Transformer & Estimator
- Optimization During Hard Parse Operation
- SQL Statement Parsing: Overview
- Structured Query Language
- Plan Generator
- Cost-Based Optimizer
- Controlling the Behavior of the Optimizer, Optimizer Features and Oracle Database Releases
- Why Do You Need an Optimizer?

Interpreting Execution Plans

- Looking Beyond Execution Plans
- Automatic Workload Repository (AWR)
- Interpreting an Execution Plan
- Using the V\$SQL_PLAN View
- What Is an Execution Plan? Where To Find Execution Plans and Viewing Execution Plans
- Reading More Complex Execution Plans and Reviewing the Execution Plan
- SQL Monitoring: Overview
- Plan Table & AUTOTRACE

Application Tracing

- End-to-End Application Tracing Challenge
- Use Enterprise Manager to Trace Services
- tkprof Output with and without Index: Example
- Location for Diagnostic Traces
- What is a Service? Use Services with Client Applications & Tracing Services
- Session Level Tracing: Example
- The trcsess Utility and SQL Trace File Contents
- Invoking the tkprof Utility and Output of the tkprof Command

Optimizer: Table and Index Operations

- Row Source Operations, Main Structures and Access Paths
- Index-Organized Tables
- Bitmap Indexes, Bitmap Operations and Bitmap Join Index
- Guidelines for Managing Indexes and Investigating Index Usage
- Using Indexes: Considering Nullable Columns
- Full Table Scan
- Indexes: Overview and B*-tree Indexes and Nulls
- Composite Indexes and Invisible Index

Optimizer Join Methods

- Hash Join and Cartesian Join
- Equijoins and Nonequijoins
- Outer Joins
- Antijoins
- Nested Loops Join
- Nested Loops Join: 11g Implementation
- Semijoins
- Sort Merge join

Optimizer: Other Operators

- Result Cache Operator
- When Are Clusters Useful?
- Filter operations and Concatenation Operations
- Count Stop Key Operator
- Min/Max and First Row Operators and Other N-Array Operations
- Sorting Operators and Buffer Sort Operator
- Inlist Iterator and View Operator
- UNION [ALL], INTERSECT, MINUS

Case Study: Star Transformation

- Star Transformation Hints
- Using Bitmap Join Indexes
- The Star Schema Model and The Snowflake Schema Model
- Bitmap Join Indexes: Join Model 1 to 4
- Star Transformation Plan Examples
- Star Transformation
- Retrieving Fact Rows from One Dimension and from All Dimensions
- Joining the Intermediate Result Set with Dimensions

Optimizer Statistics

- Locking Statistics, Export/Import Statistics and Set Statistics
- Gathering System Statistics and Statistic Preferences
- Manual Statistics Gathering
- Table, Index and Column Statistics
- Histograms, Frequency Histograms and Histogram Considerations
- Types of Optimizer Statistics
- Multicolumn Statistics and Expression Statistics Overview
- Index Clustering Factor

Using Bind Variables

- Bind Variable Peeking
- Cursor Sharing Enhancements

- The CURSOR_SHARING Parameter
- Interacting with Adaptive Cursor Sharing
- Cursor Sharing and Different Literal Values
- Forcing Cursor Sharing
- Adaptive Cursor Sharing
- Cursor Sharing and Bind Variables

Using SQL Tuning Advisor

- Tuning SQL Statements Automatically
- Database Control and SQL Tuning Advisor
- Stale or Missing Object Statistics and SQL Statement Profiling
- SQL Tuning Loop, Access Path Analysis and SQL Structure Analysis
- Application Tuning Challenges
- SQL Tuning Advisor: Overview
- Implementing Recommendations
- Plan Tuning Flow and SQL Profile Creation

Using SQL Access Advisor

- SQL Access Advisor: Overview
- SQL Access Advisor: Schedule and Review
- SQL Access Advisor: Workload Source
- SQL Access Advisor: Results
- Possible Recommendations
- SQL Access Advisor Session: Initial Options
- SQL Access Advisor: Recommendation Options
- SQL Access Advisor: Results and Implementation

Using Automatic SQL Tuning

- Automatic SQL Tuning
- Automatic SQL Tuning: Result Summary
- Configuring Automatic SQL Tuning
- SQL Tuning Loop
- Automatic SQL Tuning: Result Details
- Automatic SQL Tuning Result Details: Drilldown
- Automatic SQL Tuning Considerations
- Automatic Tuning Process

SQL Performance Management

- SQL Plan Baseline: Architecture
- Maintaining SQL Performance and SQL Plan Management: Overview
- SQL Plan Selection
- Enterprise Manager and SQL Plan Baselines
- Possible SQL Plan Manageability Scenarios
- SQL Performance Analyzer and SQL Plan Baseline Scenario

- Loading a SQL Plan Baseline Automatically and Purging SQL Management Base Policy
- Important Baseline SQL Plan Attributes