

Oracle AI Database: Performance Management and Tuning ELS

Oracle Database

DURATION

5 Days

MODULES

41 Lectures

COURSE CODE

—

Course Overview

Additional Information Role:Database Administrator Learn about Oracle's Products:Oracle Database Solutions:Database Administration

What You Will Learn

Overview

- Objectives
- What is performance management?
- Who manages performance?
- What does the DBA tune?
- Types of Tuning
- Tuning Methodology
- Effective Tuning Goals
- General Tuning Session
- Tuning a CDB
- Performance Tuning: Diagnostics
- Performance Tuning: Features and Tools
- Tuning Objectives
- Summary

Defining the Scope of Performance Issues

- Objectives
- Defining the Problem
- Limit the Scope
- Determining Tuning Priorities
- Common Tuning Problems
- Tuning Life Cycle Phases
- Tuning During the Life Cycle

- Application Design and Development
- Testing: Database Configuration
- Deployment
- Production
- Migration, Upgrade, and Environment Changes
- ADDM Tuning Session
- Performance Versus Business Requirements
- Monitoring and Tuning Tools: Overview
- Summary

Using the Time Model to Diagnose Performance Issues

- Objectives
- Time Model: Overview
- DB Time
- CPU and Wait Time Tuning Dimensions
- Time Model Statistics Hierarchy
- Time Model: Example
- Top Timed Events
- Summary

Using Statistics and Wait Events to Diagnose Performance Issues

- Objectives
- Dynamic Performance Views
- Dynamic Performance Views: Usage Examples
- Dynamic Performance Views: Considerations
- Statistic Levels
- Instance Activity and Wait Event Statistics
- System Statistic Classes
- Displaying Statistics
- Displaying SGA Statistics
- Wait Events
- Using the V\$EVENT_NAME View
- Wait Classes
- Displaying Wait Event Statistics
- Commonly Observed Wait Events
- Using the V\$SESSION_WAIT View
- Precision of System Statistics
- Summary

Using Log and Trace Files to Monitor Performance

- Objectives
- Viewing the Alert Log
- Using Alert Log Information as an Aid in Managing Performance
- Administering the DDL Log File

- Understanding the Debug Log File
- User Trace Files
- Background Processes Trace Files
- Summary

Using Enterprise Manager Cloud Control and SQL Developer to Monitor Performance

- Objectives
- Enterprise Manager: Overview
- Oracle Enterprise Manager Cloud Control Components
- Using Features of the Oracle Management Packs and Options
- Oracle SQL Developer
- SQL Developer Command Line (SQLcl)
- Summary

Using Statspack to View Performance Data

- Objectives
- Introduction to Statspack
- Statspack Scripts
- Installing Statspack
- Capturing Statspack Snapshots
- Configuring Snapshot Data Capture
- Statspack Snapshot Levels
- Statspack Baselines and Purging
- Reporting with Statspack
- Statspack Considerations
- Statspack Reports
- Reading a Statspack Report
- Statspack Report Drilldown Sections
- Report Drilldown Examples
- Load Profile Section
- Time Model Section
- Statspack and AWR
- Summary

Using Automatic Workload Repository

- Objectives
- Automatic Workload Repository: Overview
- Automatic Workload Repository Data
- Workload Repository
- AWR Administration
- AWR Snapshot Purging Policy
- Managing Snapshots with PL/SQL
- AWR Snapshot Settings
- Manual AWR Snapshots

- Managing AWR Data in a Multitenant Environment
- AWR Snapshots and ADDM in a Multitenant Architecture Database
- Generating AWR Reports
- Reading the AWR Report
- AWR Report: Multitenant Data
- Generating AWR Reports by Using SQL*Plus
- Statspack and AWR Reports
- Reading a Statspack or an AWR Report
- Compare Periods: Benefits
- Snapshots and Periods Comparisons
- Compare Periods: Results
- Compare Periods: Report
- Multitenant AWR Views
- Summary

Using Metrics and Alerts

- Objectives
- Metrics and Alerts
- Limitation of Base Statistics
- Typical Delta Tools
- Oracle Database Metrics
- Benefits of Metrics
- Viewing Metric History Information
- Viewing Metric Details
- Statistics Histograms
- Histogram Views
- Server-Generated Alerts
- Alert Usage Model
- Metric and Alert Views
- Summary

Using Baselines

- Objectives
- Comparative Performance Analysis with AWR Baselines
- Automatic Workload Repository Baselines
- AWR Baselines
- Types of Baselines
- Moving Window Baseline
- Baseline Templates
- Creating AWR Baselines
- Creating a Single AWR Baseline
- Creating a Repeating Baseline and Template
- Managing Baselines by Using the DBMS_WORKLOAD_REPOSITORY
- Package
- Generating a Baseline Template for a Single Time Period

- Creating a Repeating Baseline Template
- Baseline Views
- Performance Monitoring and Baselines
- Summary

Managing Automated Maintenance Tasks

- Objectives
- Automated Maintenance Tasks
- Maintenance Windows
- Default Maintenance Plan
- Automated Maintenance Task Priorities
- Configuring Automated Maintenance Tasks
- Summary

Using ADDM to Analyze Performance

- Objectives
- ADDM Performance Monitoring
- ADDM and Database Time
- DB Time-Graph and ADDM Methodology
- Top Performance Issues Detected
- ADDM Recommendations
- Creating a Manual ADDM Task
- ADDM Tasks in a Multitenant Architecture Database
- Changing ADDM Attributes
- Retrieving ADDM Reports by Using SQL
- Compare Period ADDM: Analysis
- Workload Compatibility
- Configuring Automatic ADDM Analysis at the PDB Level
- Using the DBMS_ADDM Package to Compare Periods
- Example: Using the DBMS_ADDM Package to Compare Periods
- Summary

Using Active Session History Data for First Fault System Analysis

- Objectives
- Active Session History: Overview
- Active Session History: Mechanics
- ASH Sampling: Example
- Accessing ASH Data
- Analyzing the ASH Data
- Using Enterprise Manager to View ASH Analysis
- Using Enterprise Manager to Generate ASH Reports
- Using the ASH Report Script to Generate a Report
- ASH Report Structure
- Determining the Source of Data

- Performing Skew Analysis
- Additional Automatic Workload Repository Views
- Summary

Using Emergency Monitoring and Real-Time ADDM to Analyze Performance Issues

- Objectives
- Emergency Monitoring: Challenges
- Emergency Monitoring: Goals
- Using Real-Time ADDM to Perform a Root Cause Analysis
- Using Real-Time ADDM
- Real-Time ADDM in the Database
- Using Real-Time ADDM
- Viewing Real-Time ADDM Results
- Summary

Overview of SQL Statement Processing

- Objectives
- SQL Statement Processing Phases
- Parsing
- SQL Cursor Storage
- Session Cursor Cache
- Cursor Usage and Parsing
- SQL Statement Processing Phases: Bind
- SQL Statement Processing Phases: Execute and Fetch
- Processing a DML Statement
- Commit Processing
- Identifying Poorly Performing SQL Statements
- Top SQL Reports
- SQL Monitoring
- Monitored SQL Execution Details
- Summary

Maintaining Indexes

- Objectives
- Creating Indexes
- Using Invisible and Unusable Indexes
- Dropping Indexes
- Reducing the Cost of SQL Operations
- Index Maintenance
- Using Advanced Index Compression
- Other Index Options
- SQL Access Advisor
- Automatic Indexing Task
- Automatic Index Task Workflow

- Automatic Indexing Task Reporting
- Automatic Indexing Views
- Summary

Maintaining Tables

- Objectives
- Reducing the Cost of SQL Operations
- Table Maintenance for Performance
- Table Reorganization Methods
- Space Management
- Extent Management
- Locally Managed Extents
- Large Extents: Considerations
- How Table Data Is Stored
- Anatomy of a Database Block
- Minimizing Block Visits
- Block Allocation
- Free Lists
- Block Space Management
- Block Space Management with Free Lists
- Automatic Segment Space Management
- Automatic Segment Space Management at Work
- Block Space Management with ASSM
- Creating an Automatic Segment Space Management Segment
- Migration and Chaining
- Guidelines for PCTFREE and PCTUSED
- Detecting Migration and Chaining
- Selecting Migrated Rows
- Eliminating Migrated Rows
- Shrinking Segments: Overview
- Shrinking Segments: Considerations
- Shrinking Segments by Using SQL
- Segment Shrink: Basic Execution
- Segment Shrink: Execution Considerations
- Data Compression
- Advanced Row Compression: Overview
- Advanced Row Compression: Concepts
- Using Advanced Row Compression
- Advanced Row Compression for DML Operations
- Advanced Index Compression
- How does Hybrid Columnar Compression work?
- Using the Compression Advisor
- Using the Compression Advisor for Indexes
- Viewing Table Compression Information
- Summary

Introduction to Query Optimizer

- Objectives
- Role of the Oracle Optimizer
- Functions of the Query Optimizer
- Selectivity
- Cardinality and Cost
- Changing Optimizer Behavior
- Setting and Viewing Optimizer Parameters
- Using Initialization Parameters to Control Optimizer Behavior
- Enabling Query Optimizer Features
- Influencing the Optimizer Approach
- Optimizing SQL Statements
- Access Paths
- Choosing an Access Path
- Summary

Understanding Execution Plans

- Objectives
- What is an execution plan?
- Methods for Viewing Execution Plans
- Uses of Execution Plans
- DBMS_XPLAN Package: Overview
- EXPLAIN PLAN Command
- EXPLAIN PLAN Command: Example
- EXPLAIN PLAN Command: Output
- Reading an Execution Plan
- Using the V\$SQL_PLAN View
- Querying V\$SQL_PLAN
- V\$SQL_PLAN_STATISTICS View
- Querying the AWR
- SQL*Plus AUTOTRACE
- Using SQL*Plus AUTOTRACE
- SQL*Plus AUTOTRACE: Statistics
- Adaptive Execution Plans
- Dynamic Plans
- Dynamic Plan: Adaptive Process
- Dynamic Plans: Example
- Continuous Adaptive Query Plans
- Automatic Re-Optimization
- Comparing Execution Plans
- Summary

Viewing Execution Plans by Using SQL Trace and TKPROF

- Objectives

- SQL Trace Facility
- How to Use the SQL Trace Facility
- Initialization Parameters
- Enabling SQL Trace
- Disabling SQL Trace
- Formatting Your Trace Files
- TKPROF Command Options
- Output of the TKPROF Command
- TKPROF Output with No Index: Example
- TKPROF Output with Index: Example
- Generating an Optimizer Trace
- Summary

Managing Optimizer Statistics

- Objectives
- Optimizer Statistics
- Types of Optimizer Statistics
- Optimizer Statistics Collection
- Dynamic Statistics
- Gathering Statistics and Setting Optimizer Statistics Preferences
- Setting Statistic Preferences
- Viewing and Managing Optimizer Statistics Preferences
- Extended Statistics
- Maintaining Optimizer Statistics
- Automated Maintenance Tasks
- Optimizer Statistics Advisor
- Optimizer Statistics Advisor Report
- Executing Optimizer Statistics Advisor Tasks
- Restoring Statistics
- Deferred Statistics Publishing: Overview
- Deferred Statistics Publishing: Example
- Managing Real-Time Statistics
- Configuring High-Frequency Automatic Optimizer Statistics Collection
- Summary

Using Automatic SQL Tuning

- Objectives
- Automatic SQL Tuning: Overview
- SQL Statement Profiling
- Plan Tuning Flow and SQL Profile Creation
- SQL Tuning Loop
- Using SQL Profiles
- Summary

Using the SQL Plan Management Feature

- Objectives
- SQL Plan Management: Overview
- SQL Plan Baseline: Architecture
- Loading SQL Plan Baselines
- Loading SQL Plan Baselines from AWR
- Evolving SQL Plan Baselines
- Important SQL Plan Baseline Attributes
- SQL Plan Selection
- Possible SQL Plan Manageability Scenarios
- SQL Performance Analyzer and SQL Plan Baseline Scenario
- Loading a SQL Plan Baseline Automatically
- Purging SQL Management Base Policy
- Enterprise Manager and SQL Plan Baselines
- Automatic SQL Plan Management
- SPM Evolve Advisor
- Summary

Overview of the SQL Advisors

- Objectives
- SQL Tuning Process
- SQL Tuning Advisor: Overview
- SQL Access Advisor: Overview
- SQL Performance Analyzer: Overview
- Summary

Using the SQL Tuning Advisor

- Objectives
- SQL Tuning Advisor: Overview
- SQL Tuning Advisor Architecture
- Automatic Tuning Optimizer
- Using the SQL Tuning Advisor
- SQL Tuning Advisor Options
- SQL Tuning Advisor Recommendations
- Alternative Execution Plans
- Summary

Using the SQL Access Advisor

- Objectives
- SQL Access Advisor: Overview
- Using the SQL Access Advisor
- Viewing Recommendations
- Viewing Recommendation Details

- Practice Overview
- Summary

Overview of Real Application Testing Components

- Objectives
- Real Application Testing: Overview
- Real Application Testing: Use Cases
- Summary

Using SQL Performance Analyzer to Determine the Impact of Changes

- Objectives
- SQL Performance Analyzer: Process
- Steps 6–7: Comparing/Analyzing Performance and Tuning Regressed SQL
- Capturing the SQL Workload
- Creating a SQL Performance Analyzer Task
- SQL Performance Analyzer Task Page
- SQL Performance Analyzer Comparison Report
- SQL Performance Analyzer: PL/SQL Example
- Tuning Regressed SQL Statements
- SQL Performance Analyzer: Data Dictionary Views
- Summary

Using Database Replay to Test System Performance

- Objectives
- Using Database Replay
- The Big Picture
- System Architecture: Capture
- System Architecture: Processing the Workload
- System Architecture: Replay
- Database Replay Workflow in Enterprise Manager
- Accessing Database Replay in Enterprise Manager
- Capture Considerations
- Replay Considerations: Preparation
- Replay Considerations
- Replay Customized Options
- Replay Analysis
- Database Replay Packages
- Data Dictionary Views: Database Replay
- Database Replay: PL/SQL Example
- Calibrating Replay Clients
- Capturing and Replaying in a CDB and PDBs
- Reporting
- Summary

Implementing Real-Time Database Operation Monitoring

- Objectives
- Overview
- Use Cases
- Defining a DB Operation
- Scope of a Composite DB Operation
- Database Operation Concepts
- Identifying a Database Operation
- Enabling Monitoring of Database Operations
- Identifying, Starting, and Completing a Database Operation
- Monitoring Database Operations in Sessions
- Monitoring the Progress of a Database Operation
- Monitoring SQL Operation Details
- Database Operation View: V\$SQL_MONITOR
- Database Operation Views
- Reporting Database Operations by Using Functions
- Database Operation Tuning
- Summary

Using Services to Monitor Applications

- Objectives
- What is a service?
- Service Attributes
- Service Types
- Creating Services
- Using the DBMS_SERVICE Package to Manage Services
- Where are services used?
- Using Services with Client Applications
- Using Services with the Resource Manager
- Using Enterprise Manager to Manage Consumer Group Mappings
- Services and the Resource Manager: Example
- Using Enterprise Manager to Create a Job Class
- Using Enterprise Manager to Create a Job
- Services and the Scheduler: Example
- Using Services with Metric Thresholds
- Using Enterprise Manager to Change Service Thresholds
- Services and Metric Thresholds: Example
- Service Aggregation and Tracing
- Services Statistics in Performance Home Page
- Top Services Statistics Using ASH Analytics Performance Page
- Service Aggregation Configuration
- Service Aggregation: Example
- Client Identifier Aggregation and Tracing
- Using the TRCSESS Utility

- Service Performance Views
- Summary

Overview of Memory Structures

- Objectives
- Managing Memory Caches and Structures
- Guidelines for Efficient Memory Usage
- Unified Memory
- Summary

Managing Shared Pool Performance

- Objectives
- Shared Pool Architecture
- Shared Pool Operation
- Library Cache
- Latch and Mutex
- Latch and Mutex: Views and Statistics
- Diagnostic Tools for Tuning the Shared Pool
- AWR/Statspack Indicators
- Top Timed Events
- Time Model
- Load Profile
- Instance Efficiencies
- Library Cache Activity
- Avoid Hard Parses
- Are cursors being shared?
- Candidate Cursors for Sharing
- Sharing Cursors
- Adaptive Cursor Sharing: Example
- Adaptive Cursor Sharing Views
- Interacting with Adaptive Cursor Sharing
- Reducing the Cost of Soft Parses
- Sizing the Shared Pool
- Shared Pool Advisory
- Shared Pool Advisory in an AWR Report
- Shared Pool Advisor
- Avoiding Fragmentation
- Large Memory Requirements
- Tuning the Shared Pool Reserved Pool
- Keeping Large Objects
- Data Dictionary Cache
- Dictionary Cache Misses
- SQL Query Result Cache: Overview
- Managing the SQL Query Result Cache
- Using the RESULT_CACHE Hint

- Using Table Annotation to Control Result Caching
- Using the DBMS_RESULT_CACHE Package
- Viewing SQL Result Cache Dictionary Information
- SQL Query Result Cache: Considerations
- Summary

Managing Buffer Cache Performance

- Objectives
- Buffer Cache: Highlights
- Database Buffers
- Buffer Hash Table for Lookups
- Working Sets
- Tuning Goals and Techniques
- Symptoms of a Buffer Cache Issue
- Cache Buffer Chains Latch Contention
- Finding Hot Segments
- Buffer Busy Waits
- Buffer Cache Hit Ratio
- Buffer Cache Hit Ratio is Not Everything
- Interpreting Buffer Cache Hit Ratio
- Read Waits
- Free Buffer Waits
- Solutions for Buffer Cache Issues
- Sizing the Buffer Cache
- Buffer Cache Size Parameters
- Dynamic Buffer Cache Advisory Parameter
- Buffer Cache Advisory View
- Using the V\$DB_CACHE_ADVICE View
- Using the Buffer Cache Advisor
- Caching Tables
- Automatic Big Table Caching
- Configuring Automatic Big Table Caching
- Using Automatic Big Table Caching
- Monitoring Automatic Big Table Caching
- Memoptimized Rowstore
- In-Memory Hash Index
- Multiple Buffer Pools
- Enabling Multiple Buffer Pools
- Calculating the Hit Ratio for Multiple Pools
- Multiple Block Sizes
- Multiple Database Writers
- Multiple I/O Slaves
- Using Multiple Writers and I/O Slaves
- Private Pool for I/O-Intensive Operations
- Automatically Tuned Multiblock Reads

- Database Smart Flash Cache Overview
- Using Database Smart Flash Cache
- Database Smart Flash Cache Architecture: Overview
- Configuring Database Smart Flash Cache
- Sizing Database Smart Flash Cache
- Enabling and Disabling Flash Devices
- Specifying Database Smart Flash Cache for a Table
- Full Database In-Memory Caching
- Setting Up Force Full Database Caching
- Monitoring Full Database In-Memory Caching
- Flushing the Buffer Cache (for Testing Only)
- Summary
- Practice Overview

Managing PGA and Temporary Space Performance

- Objectives
- SQL Memory Usage
- Performance Impact
- Automatic PGA Memory
- SQL Memory Manager
- Configuring Automatic PGA Memory
- Setting PGA_AGGREGATE_TARGET Initially
- Limiting the Size of the Program Global Area
- Managing the PGA for PDBs
- Monitoring SQL Memory Usage
- Monitoring SQL Memory Usage: Examples
- Tuning SQL Memory Usage
- PGA Target Advice Statistics
- PGA Target Advice Histograms
- Automatic PGA and Enterprise Manager
- Automatic PGA and AWR Reports
- Temporary Tablespace Management: Overview
- Temporary Tablespace: Locally Managed
- Configuring Temporary Tablespace
- Temporary Tablespace Group: Overview
- Temporary Tablespace Group: Benefits
- Creating Temporary Tablespace Groups
- Maintaining Temporary Tablespace Groups
- Viewing Tablespace Groups
- Monitoring Temporary Tablespace
- Shrinking a Temporary Tablespace
- Using the Tablespace Option When Creating a Temporary Table
- Summary

Configuring the Large Pool

- Objectives
- Large Pool Overview
- Tuning the Large Pool
- Summary

Using Automatic Shared Memory Management

- Objectives
- Oracle Database Architecture
- Granules
- Automatic Shared Memory Management: Overview
- SGA Sizing Parameters: Overview
- Dynamic SGA Transfer Modes
- Memory Broker Architecture
- Manually Resizing Dynamic SGA Parameters
- Behavior of Auto-Tuned SGA Parameters
- Behavior of Manually Tuned SGA Components
- Using the V\$SYSTEM_PARAMETER View
- Resizing SGA_TARGET
- Disabling Automatic Shared Memory Management
- Using the SGA Advisor
- Monitoring ASMM
- Managing SGA for PDBs
- Summary

Introduction to In-Memory Column Store

- Objectives
- Database In-Memory Feature Set
- Goals of In-Memory Column Store
- Benefits
- Overview
- Row Store Versus Column Store: 2D Vision
- In-Memory Column Unit
- Compare: In-Memory Column Store Cache and Buffer Cache
- Dual Format In-Memory
- Indexes Issues
- Process
- In-Memory Column Store: Dual Format of Segments in SGA
- Using OEM to Manage In-Memory Pool
- Summary

Configuring the In-Memory Column Store Feature

- Objectives

- Deploying IM Column Store
- Using OEM to Manage In-Memory Pool
- Deploying IM Column Store: Objects Setting
- Deploying IM Column Store: Columns Setting
- Defining IM Column Store Compression
- In-Memory Advisor
- IM Advisor or Compression Advisor?
- Computing Compression Ratio
- IM FastStart
- Automatic In-Memory: Overview
- AIM Action
- Configuring Automatic In-Memory
- Diagnostic Views
- Summary

Using the In-Memory Column Store Feature to Improve SQL Performance

- Objectives
- Query Benefits
- Testing and Comparing Query Performance
- Queries on In-Memory Tables: Simple Predicate
- MINMAX Pruning Statistics
- IM Column Store Statistics
- Execution Plan: TABLE ACCESS IN MEMORY FULL
- Queries on In-Memory Tables: Join
- Execution Plan: JOIN FILTER CREATE / USE
- Queries on In-Memory Tables: Join Groups
- Population of Expressions and Virtual Columns Results
- In-Memory Expression Unit (IMEU)
- Populating In-Memory Expression Results
- Populating In-Memory Expression Results Within a Window
- Waiting for In-Memory Segments to Be Populated
- Views
- Summary

Using In-Memory Column Store with Oracle Database Features

- Objectives
- Interaction with Other Products
- Optimizer
- IM Column Store and Real Application Clusters
- IM Column Store and Data Pump
- Data Pump TRANSFORM Names
- Automatic Data Optimization Interaction
- Managing Heat Map and Automatic Data Optimization Policies
- Creating ADO In-Memory Policies
- Summary