Oracle AI Database: New Features for Administrators LVC

Oracle Database

DURATION	MODULES	COURSE CODE
3 Days	**19 Lectures**	**—**

Course Overview

This course covers the latest additions and improvements in Oracle AI Database. It equips Oracle Database Administrators with knowledge about enhanced tools for better management, focusing on performance, architecture, security, and database sharding. The topics include updates to Blockchain Tables, Data Migration with RMAN, Wide Tables, and Lock-Free Reservations.

What You Will Learn

- I Course Overview
- Presenters I-2
- Prerequisites I-4
- Learning Outcome I-5
- Course Objectives I-6

Oracle AI Vector Search

- Objectives
- Oracle AI Vector Search Benefits
- Benefits of Oracle AI Vector Search
- VECTOR Data Type
- Examples
- The Complete Workflow
- Vector Embeddings
- Vector Embedding Models
- Import Embedding Models
- Basic Queries and Similarity Search
- Basic Queries
- Basic Queries: Comparison Operations
- Similarity Search
- Exact Similarity Search
- Vector Distance Metrics
- Vector Distance Metric Example: Exact Similarity Search

- Euclidean Metric
- Euclidean Squared Distance Metric Example
- Approximate Similarity Search
- Approximate Similarity Search or Exact Similarity Search?
- Approximate Similarity Search
- Approximate Similarity Search: HNSW
- Approximate Similarity Search: IVF

True Cache

- Objectives
- About True Cache
- True Cache Benefits
- True Cache: High-Level View
- True Cache: Good to Know
- True Cache Application Usage
- Concurrency Control
- Configuring and Deploying True Cache
- Configuring True Cache Overview
- Uniform Configuration
- Partition Configuration COLOCATION_TAG
- Partition Configuration Multiple Services
- True Cache Configuration Process Overview
- True Cache Best Practices
- Best Practices for Maximum Availability Architecture (MAA)
- Configuring and Deploying True Cache: With DBCA
- DBCA Prerequisites
- Configuring True Cache with DBCA
- Configuring True Cache Database Application Services
- Configuring and Deploying True Cache: Manually
- Manually Configuring True Cache: Prerequisites
- Manually Configuring True Cache: Simplified Steps
- Configuring Oracle True Cache: tnsnames.ora
- Configuring Oracle True Cache: tnanames.ora
- Configuring Oracle True Cache: Local Listener
- Configuring Oracle True Cache: Password File
- Configuring Oracle True Cache: Prepare PFILE
- Configuring Oracle True Cache: Initial True Cache Startup
- Configuring Oracle True Cache: Create True Cache Database Instance
- Configuring and Deploying True Cache: Application Services
- Configuring True Cache Database Application Services
- Creating Database Application Services on the Primary Database
- Verifying Database Application Services Are Created
- Starting Corresponding Database Application Services
- Configuring and Deploying True Cache: RAC Primary Database
- Deploying True Cache for an Oracle RAC Primary Database

- Configuring and Deploying True Cache: Verify Configuration
- Configuring Oracle True Cache: Verify Status
- Configuring Oracle True Cache: Verify Listener
- Configuring Oracle True Cache: Verify Application Services
- Configuring and Deploying True Cache: Enable DML Redirection
- Configuring Oracle True Cache: Enable DML Redirection
- Configuring and Deploying True Cache: Deploy in Containers
- Deploying True Cache in Containers
- Managing True Cache: Shutdown and Start
- Shutting Down and Starting True Cache
- Managing True Cache: Deleting True Cache
- Deleting True Cache: DBCA
- Deleting True Cache: Manually
- Monitoring True Cache
- Monitoring True Cache: V$TRUE_CACHE View
- Monitoring True Cache: V$TRUE_CACHE View Columns
- Monitoring True Cache: Automatic Workload Repository
- Using Oracle True Cache in Your Applications
- True Cache Application: JDBC
- Best Practices for Load Balancing in a Uniform Configuration
- Complementary Caching Features
- Test Application
- Summary

Sharding New Features

- Objectives
- Database Sharding
- Oracle Database Sharding
- Oracle Database 23ai Sharding New Features
- Sharding Native Replication (RAFT-Based)
- Directory-Based Sharding
- Coordinated Backup and Restore for Sharded Databases
- Centralized Backup Recovery
- Automatic Bulk Data Move on Sharding Key Update
- Split/Move of a Partitionset with Bulk Data Movement
- Summary

Automatic SQL Plan Management

- Objectives
- Automatic SQL Plan Management: Overview
- SPM Evolve Advisor
- Benefits
- Problems It Solves
- When to Use
- How to Use It

- Summary

Automatic Transaction Rollback

- Objectives
- Overview
- Using Automatic Transaction Rollback
- Automatic Transaction Rollback
- Using Automatic Transaction Rollback
- Monitoring
- Benefits
- Summary

Automatic Transaction Quarantine

- Objectives
- Overview
- Automatic Transaction Quarantine: Workflow
- Monitoring Quarantined Transactions
- Resolving Quarantined Transactions
- Resolving Quarantined Transactions Examples
- Resolving Quarantined Transactions: Examples
- Dropping Quarantined Transactions
- Quarantined Transaction Escalation
- Summary

Flashback Log Placement

- Objectives
- Flashback Log Overview
- Flashback Log Parameters
- Flashback Log Management
- Summary

Simplified Database Migration Across Platforms Using RMAN Backups

- Objectives
- Overview
- Methods of Transporting with Recovery Catalog
- Methods of Transporting in NOCATALOG Mode
- Methods of Transporting Over Network Link
- Prerequisites Generic
- Prerequisites Specific
- Recovery Catalog Methods for PDBs
- Quickly Transport a PDB with Recovery Catalog
- Transport a PDB by Using a Pre-existing Backup with Recovery Catalog
- Transport a PDB Using Multiple Incremental Backups with Recovery Catalog
- Recovery Catalog Methods for Tablespaces

- Recovery Catalog Method for Tablespaces
- Quickly Transport a Tablespace with Recovery Catalog
- Transport a Tablespace by Using a Pre-existing Backup with Recovery
- Catalog
- Transport a Tablespace Using Multiple Incremental Backups with Recovery
- Catalog
- Transport a PDB Using Multiple Incremental Backups with Recovery Catalog
- Methods of Transporting in NOCATALOG Mode
- NOCATALOG Mode for PDBs
- Quickly Transport a PDB NOCATLOG
- Transport a PDB Using a Pre-existing Backup NOCATLOG
- Transport a PDB Using Multiple Incremental Backups NOCATLOG
- NOCATALOG Mode for Tablespaces
- Quickly Transport a Tablespace NOCATALOG
- Transport a Tablespace by Using a Pre-existing Backup NOCATALOG
- Transport a Tablespace Using Multiple Incremental Backups NOCATALOG
- Methods of Transporting Over Network Link
- Transporting PDBs Over Network Link
- Quickly Transport a PDB Over Network Link
- Transporting PDBs by Restoring Backups Incrementally Over the Network
- Methods of Transporting Over Network Link
- Transporting Tablespaces Over Network Link
- Quickly Transport a Tablespace Over Network Link
- Transporting Tablespaces by Restoring Backups Incrementally Over the
- Network
- Summary

Lock-Free Reservations

- Objectives
- Optimistic vs Lock Free
- Optimistic Approach
- Lock-Free Approach
- Frequently Modified Values
- Lock-Free Reservations
- Lock-Free Reservation Features
- Lock-Free Reservation Benefits
- Functionality Requirements
- Create, Modify, Delete
- Lock-Free Journal Table
- Journal Table Restrictions
- Lock-Free Reservations
- Views
- Summary

Wide Columns

- Objectives
- Prior Challenges
- Benefits
- Increase Column Limits
- Performance
- Summary

Improved Performance Hybrid Columnar Compression (HCC)

- Objectives
- HCC Pre 23ai
- HCC in 23ai
- Database Compatibility
- Performance
- Summary

Unrestricted Parallel DMLs

- Objectives
- Overview Parallel DML
- 23ai Parallel DML Improvements
- Benefits
- Summary

Unrestricted Direct Loads

- Objectives
- Direct Loads Pre 23ai
- Direct Loads 23ai
- Benefits
- Summary

Accelerate SecureFiles LOB Write Performance

- Objectives
- Overview
- Benefits
- Summary

Fast Inserts Enhancements

- Objectives
- Fast Ingest Overview (aka Deferred Inserts)
- Fast Ingest Overview
- Using Fast Ingest
- Enhancements

- Benefits
- Management
- Summary

Blockchain Tables

- Objectives
- Blockchain Table Overview
- Blockchain Table Row Versions
- Blockchain Table User Chains
- Blockchain Table Delegate Signer
- Blockchain Table Countersignature
- Summary

Security New Features

- Objectives
- Database Security
- Pluggable Database Hybrid Read-Only Mode
- Read-Only User and Session
- New Developer Role
- Simplified Schema Privileges
- Simplified Schema Privileges: Example
- Audit Object Actions at the Column Level
- Ability to Control Authorizations for Unified Auditing with Database Vault
- Integration of Microsoft Azure Active Directory with Additional Oracle Database
- Environments
- Summary

Data Storage Improvements

- Objectives
- SecureFiles
- Traditional SecureFiles Shrinking
- Automatic SecureFiles Shrinking: Advantages
- High-Level View
- Controlling Automatic SecureFiles Shrink
- Automatic Storage Compression
- Automatic Storage Compression and Automatic Clustering
- Automatic Storage Compression: Advantages
- Controlling Automatic Storage Compression
- Summary

Time and Date Handling Changes

- Objectives
- Enhanced Time Zone Data Upgrade
- Time Zone Definitions Change: Considerations

- Optimizing Time Zone File Upgrade
- Benefits of Enhanced Time Zone Update
- Tables with TIMESTAMP WITH TIME ZONE Data
- New Parameter in the init.ora File
- Configuring Database Parameters
- SYSDATE and SYSTIMESTAMP
- Date and Time in Oracle Databases
- Database Time Versus OS System Time
- Database-Specific Time in Oracle Database 23ai
- Setting Database-Specific Time: Example
- Summary