

Oracle Database 12c: OCM Exam Preparation Workshop Ed 1

Duration: 5 Days

What you will learn

The Oracle Database 12c: OCM Exam Preparation Workshop is designed for those candidates who are preparing to take the Oracle Database 12c: Oracle Certified Master (OCM) Exam. This workshop has extensive hands-on practices to enable students to:

- Create a Multitenant Container Database and Pluggable Databases
- Perform Backup, Recovery and Flashback of CDB and PDBs
- Load and unload data
- Implement partitioning
- Configure Dataguard
- Configure clusterware
- Install RAC

Audience

- Data Center Manager
- Data Warehouse Administrator
- Data Warehouse Developer
- Database Administrators
- Database Designers
- Systems Architects
- Technical Administrator
- Technical Consultant

Course Objectives

- Describe the multitenant architecture
- Establish connections to CDB / PDB
- Implement fine-grained auditing (FGA)
- Describe operations that can be parallelized
- Describe ways to move data
- Describe the architecture, uses, and advantages of partitioning
- Describe how to optimize a star query using star transformation
- Use the SQL Access Advisor

Explain the goals, benefits, and architecture of the inmemory column store

Explain the differences between physical and logical standby databases

Use real-time query to access data on a physical standby database

Create a snapshot standby database to meet the requirement for a temporary, updatable snapshot of a physical standby database

Explain the principles and purposes of clusters

Explain and apply Automatic Storage Management (ASM) initialization parameters

Describe the benefits of Oracle RAC

Modify initialization parameters in a RAC environment

Course Topics

Basics of Multitenant Container Database and Pluggable Databases

New Multitenant Architecture: Benefits

Multitenant Container Database

Separating SYSTEM and User Data

Provisioning a Pluggable Database

Interacting Within Multitenant Container Database

Multitenant Container Database Architecture

Common and Local Users

Common and Local Privileges and Roles

Managing Tablespaces and Users in CDB and PDBs

Tablespaces in PDBs

Creating Permanent Tablespaces in a CDB

Creating Local Temporary Tablespaces

Tablespace Encryption: Advantages

Creating an Encrypted Tablespace

Common and Local Schemas/Users

Granting and Revoking Privileges

Granting Common or Local Privileges/Roles to Roles

Basics of Parallel Execution

Parallelizable Operations

Previewing Execution Plans

Understanding Parallel Execution: Query Coordinator (QC) and Parallel (PX) Servers

Parallel Execution Plans: Identifying the Producers and Consumers

Parallel Execution Plans: Understanding Granules

Example PX Parameters

PX Parameters for Auto DOP

Parameters for PX Messaging: Shared and Large Pools

In-Memory Parallel Execution & Multi-Threaded architecture

- Direct Reads Versus Buffer Cache Reads
- Parallel Execution and the Buffer Cache
- In-Memory Parallel Execution
- When In-Memory Parallel Execution Works
- Controlling In-Memory Parallel Execution
- Enhance In-Memory PX Using Server Pools
- Enhance In-Memory PX: Automatic Big Table Caching
- Using PARALLEL_FORCE_LOCAL Parameter

Moving Data

- Oracle Data Pump: Overview
- Directory Objects for Data Pump
- Data Pump Export and Import Clients: Overview
- External Tables
- Data Pump Import: Transformations
- SQL*Loader: Overview
- Loading Methods
- Transportable Tablespaces

Working with Indexes

- Normal B*-tree Indexes
- Index Scans
- Index Range Scan
- Index Range Scan: Function-Based
- Index Fast Full Scan
- B*-tree Indexes and Nulls
- Index-Organized Tables
- Bitmap Indexes

Partitioning Concepts

- Partitioned Tables and Indexes
- Partitioning Strategies: Single-Level Partitioning
- Oracle Partitioning History
- Partition Key Extensions
- Benefits of Partitioning: Table Availability
- Manageability: Relocate Table Data
- Benefits of Partitioning: Performance Considerations
- Verifying Partition Use

Dataware House Tuning and Partitioning Workloads

- Characteristics of a Data Warehouse
- OLTP Systems Versus Data Warehouses
- Data Warehouse Architectures: Basic Data Warehouse with Staging Area
- Data Warehousing Objects
- Optimizing Star Queries: Star Transformation
- Execution Plan Without Star Transformation
- Star Transformation Considerations
- Retrieving Fact Rows from All Dimensions: Phase 1

Using Flashback Technologies

- Flashback Technologies Error Detection and Correction
- Flashback Technology

- Guaranteeing Undo Retention
- Flashback Query
- Flashback Table: Overview
- Flashback Table: Considerations
- Flashback Transaction Query
- Flashing Back a Transaction

Influencing the Optimizer

- Functions of the Query Optimizer
- Cardinality and Cost
- Changing Optimizer Behavior
- Optimizer Statistics
- Optimizer Parameters
- Enabling Query Optimizer Features
- Dynamic Plans
- Cardinality Feedback: Monitoring Query Executions

SQL Performance Management and SQL Data Compression

- Maintaining SQL Performance
- Maintaining Optimizer Statistics
- Automated Maintenance Tasks
- Setting Statistic Preferences
- Restoring Statistics
- Deferred Statistics Publishing: Example
- Automatic SQL Tuning: Overview
- Plan Tuning Flow and SQL Profile Creation

In-Memory Column Store

- Goals of In-Memory Column Store
- Store Versus Column Store: 2D Vision
- In-Memory Column Unit
- Dual Format In Memory
- Deploying IM Column Store
- Objects Candidates for IM Column Store
- Defining IM Column Store Priority
- Default In-Memory Setting

Introduction to Oracle Data Guard

- What Is Oracle Data Guard?
- Types of Standby Databases
- Role Transitions: Switchover and Failover
- Choosing an Interface for Administering a Data Guard Configuration
- Primary Database Processes
- Standby Database Processes
- Physical Standby Database: Redo Apply Architecture
- Automatic Gap Detection and Resolution

Using Oracle Active Data Guard

- Using Real-Time Query
- Disabling Real-Time Query
- Checking the Standby's Open Mode
- Monitoring Apply Lag: V\$DATAGUARD_STATS

- Allowed Staleness of Standby Query Data
- Setting STANDBY_MAX_DATA_DELAY by Using an AFTER LOGON Trigger
- Forcing Redo Apply Synchronization
- Creating an AFTER LOGON Trigger for Synchronization

Creating a Logical Standby Database

- Benefits of Implementing a Logical Standby Database
- Logical Standby Database: SQL Apply Architecture
- Preparing to Create a Logical Standby Database
- Unsupported Data Types
- Checking for Tables with Unsupported Data Types
- Unsupported PL/SQL-Supplied Packages
- Adding a Disabled Primary Key RELY Constraint
- Creating a Logical Standby Database by Using SQL Commands

Enabling Fast-Start Failover

- When Does Fast-Start Failover Occur?
- Installing the Observer Software
- Fast-Start Failover Prerequisites
- Configuring Fast-Start Failover
- Step 1: Specify the Target Standby Database
- Setting the Lag-Time Limit
- Configuring the Primary Database to Shut Down Automatically
- Configuring Automatic Reinstatement of the Primary Database

Introduction to Clusterware

- Clusterware Architecture and Cluster Services
- Goals for Oracle Clusterware
- Oracle Clusterware Fencing
- Cluster Time Synchronization
- Oracle Clusterware Networking
- Grid Naming Service (GNS)
- Grid Naming Service Configuration Options
- Single-Client Access Name

ASM Overview

- What Is Oracle ASM?
- ASM Features and Benefits
- ASM Instance Designs: Clustered ASM for Clustered Databases
- ASM Components: Software
- ASM Components: ASM Instance Primary Processes
- ASM Components: Configuration Files
- ASM Disk Group: Failure Groups
- ASM Files: Extents and Striping

Flex Clusters

- Flex Clusters: Overview
- Flex Cluster Scalability
- Leaf Node Characteristics
- Cluster Mode: Overview
- Configuring the Cluster Mode
- Configuring Miss Count for Leaf Nodes

Configuring a Flex Cluster with OUI: Configuring GNS
Configuring a Flex Cluster with OUI: Selecting the Node Type

Flex ASM

Flex ASM: Overview
ASM Instance Changes
Configuring Flex ASM on a Standard Cluster
Configuring Flex ASM on a Flex Cluster
Stopping, Starting, and Relocating Flex ASM Instances
Setting the Cardinality for Flex ASM Instances
Monitoring Flex ASM Connections
Relocating an ASM Client

RAC Databases Overview and Architecture

Typical Oracle RAC Architecture
Oracle RAC One Node
Cluster-Aware Storage Solutions
Oracle RAC and Network Connectivity
Clusters and Scalability
Levels of Scalability
Speedup/Scaleup and Workloads
Necessity of Global Resources

Upgrading and Patching Oracle RAC

Patch and Patch Set: Overview
Configuring the Software Library
Reduced Down-Time Patching for Cluster Environments
Out-of-Place Database Upgrades
OPatch: Overview
Before Patching with OPatch
OPatch Automation: Examples
OPatch Log and Trace Files