Oracle Database 12c R2: New Features for Administrators Part 1 Ed 1

Duration: 5 Days

What you will learn

Throughout the lessons of the Oracle Database 12c R2: New Features for Administrators Part 1 course constituted by three modules; the Using SQL New Features and Tools, the Understanding DB Architectures, and the Diagnosing Database Problems, students receive a good knowledge of the Oracle Database 12c Release 2 new and enhanced features in different areas of database administration, mainly in on-premises Multitenant databases.

Learn To:

Use the new SQLcl utility and gain information on SQL*Plus and SQL enhancements.
Understand the new and enhanced features of Oracle Database 12c Release 2 (12.2.0.1) amongst different areas such as the new Multitenant architecture and the Sharded database architecture.
Diagnose database issues, handling Automatic Diagnostic Repository (ADR) file management and using the Trace File Analyzer (TFA) tool.

 Benefits To You

You will benefit from learning the following lessons:

Using the SQL New Features and Tools module covers the usage of the new SQLcl utility and the new features in SQL*Plus and SQL.

Most of the Understanding DB Architectures module lessons cover the new Multitenant architecture. The concept of regular and application containers of a Multitenant container database includes new administrative tasks such as creating, managing, securing, backing up, recovering and upgrading regular and application containers, as well as migrating an 11g database to a 12c multitenant container database. The lessons also cover enhanced features of Resource Manager to help you manage the resources within a Multitenant container database. You will discover other enhancements supported in a Multitenant container database like heat map and Automatic Data Optimization (ADO) support, Database Vault and encryption. The second part of the module gives an overview of database sharding. It describes the challenges and benefits of a sharded database, the sharded database architecture and how to configure a sharded database.

The lesson of the Diagnosing Database Problems module covers enhancements in ADR space management and explain what can be completed with the Trace File Analyzer (TFA) tool in terms of database diagnosis.

Audience

Database Administrators
End Users
System Administrator

Course Objectives
Create and manage pluggable databases

Describe the challenges and benefits of a sharded database

Handle the ADR automatic file space

Diagnose database issues by using the Trace File Analyzer collector

Use the new SQL*Plus history commands and SQL enhancements such as long identifiers

Use the SQLcl new utility as the new SQL command line interface

Create and manage a multitenant container database

Course Topics

Introduction
Global objectives of the course
Lessons grouped by modules
Using SQL new features and tools
Understanding DB architectures
Diagnosing database problems
Schedule of the week

Using SQL New Features and SQLcl
Use new 128 bytes identifier length for database objects
Increase length limits of data types
Use SQL row-limiting clause

Describe the support for invisible and hidden columns
Use the new VALIDATE_CONVERSION function
Recall SQL*Plus commands in the same session from history
Describe the new SQLcl utility

Understanding CDB Basics
Challenges
New Multitenant Architecture: Benefits
Non-CDB Architecture vs Oracle Multitenant Container Database
Configurations
A pristine installation
SYSTEM objects in the USER container
Provisioning a pluggable database
CDB root and pluggable database containers

Creating CDB and Regular PDBs
Create a CDB using new clauses
What's new in CDB after CDB creation
Data dictionary views
What to do in CDB after CDB creation
ADR
Provisioning new PDBs: overview
Tools

Creating Application PDBs and Installing Applications
Regular PDBs vs application PDBs
PDBs and applications
Application containers
Types of containers
Create and manage an application container
Install applications
Patch and upgrade applications
Application common objects

Creating PDBs
Configure and use locla UNDO mode
Cloning regular and application containers PDBs
Plugging unplugged regular and application PDBs into CDB
Cross-Platform transportable PDB
Plugging or cloning a non-CDB into a CDB
Perform hot cloning and relocation
Converting regular PDBs to application root or PDBs
Plugging unplugged PDBs with encrypted data

Managing CDB and PDBs
Managing CDB and PDBs
Switching Connection
Creating and renaming Services
Starting Up a CDB Instance
Mounting a CDB
Opening CDBs and PDBs
Changing the different modes and settings of PDBs

Evaluating the impact of parameter value changes

Managing Storage
Creating Permanent Tablespaces in a CDB
Objects in Tablespaces
Tablespaces Created During PDB Creation
Defining Default Permanent Tablespaces
Temporary Tablespaces
UNDO Tablespaces

Managing Security
Creating common users, roles and profiles in CDB and PDBs
Granting privileges commonly in CDB and PDBs
Common objects in Application PDBs and operations on Data-Linked objects
Enabling Common Users to Access Data in PDBs
Managing PDB lockdown profiles
Auditing users in CDB and PDBs
Protecting data with Database Vault policies in CDB and PDBs
Encrypting data in PDBs

Backing up, Recoverying and Flashing Back
New syntax and clauses in RMAN
CDB and PDB backups
Using RMAN backup to plug an unplugged PDB
Instance failure and instance recovery
PDB tempfile, essential and non-SYSTEM tablespaces recovery
PDB point-in-time recovery
Duplicating PDBs
CDB and PDB flashback

Managing Performance
Tuning a CDB
Sizing the CDB
Managing SGA and PGA for PDBs
Monitoring PDB memory usage
AWR and ADDM behaviour at CDB and PDB levels
PDB-Level snapshot views
AWR report
Controlling the Degree of Parallelism of Queries

Managing Resources Allocation
Allocating resources in the CDB and PDBs
Managing resources between PDBs
CDB Resource plan basics: Limits + cpu_count init parm per PDB
Controlling PDB IO rate limit
Managing resources within a PDB
Creating and setting PDB performance profiles

Moving and Migrating Data
Using Oracle Data Pump with PDBs
Exporting non-CDB data and importing non-CDB data into PDB
Exporting and importing between PDBs

Exporting from PDB and importing into a non-CDB
Full Transportable export/import
Transporting a database over the network
Using SQL*Loader with PDBs

Performing Miscellaneous Operations
Using catcon.pl utility to install/remove options from CDB/PDB
Using Xstreams with a CDB and PDB
Creating a standby of a CDB
Scheduling operations in a PDB
Jobs coordinator and resources
Mining statements of a PDB using LogMiner

Understanding Database Sharding
What Is Database Sharding?
Benefits of Sharding
Advantages of Oracle Sharding over NoSQL
Application Considerations for Sharding
Components of Database Sharding
Complete Deployment of a System-Managed SDB
Creating Sharded Tables
Sharded Table Family

Diagnosing Database Problems
Automatic Diagnostic Repository
New ADRCI Command
ADR Retention
Network Performance
Trace File Analyzer (TFA) Collector process and repository
TFA Collector utility
Tracing Data Pump
MVs refreshed statistics history