



# Oracle AI Database: High Availability New Features

Student Guide  
D1106068GC10

**Copyright © 2025, Oracle and/or its affiliates.**

## **Disclaimer**

This document contains proprietary information and is protected by copyright and other intellectual property laws. The document may not be modified or altered in any way. Except where your use constitutes "fair use" under copyright law, you may not use, share, download, upload, copy, print, display, perform, reproduce, publish, license, post, transmit, or distribute this document in whole or in part without the express authorization of Oracle.

The information contained in this document is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

## **Restricted Rights Notice**

If this documentation is delivered to the United States Government or anyone using the documentation on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

## **Trademark Notice**

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

## **Third-Party Content, Products, and Services Disclaimer**

This documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

1011032025

# Contents

## ----- Important Note -----

**Oracle AI Database 26ai** has replaced **Oracle Database 23ai**. This change was announced at Oracle AI World in October 2025. The architecture, concepts, and features presented in this course remain fully relevant to 26ai.

-----

### 1 Oracle Clusterware

- Objectives 1-2
- Improve Monitoring and Auto Correct Resource States 1-4
- Typical Problems Applicable and Remedies 1-5
- Primary CRSD Hang: Example 1-6
- Unmounted OCR Disk Group: Example 1-7
- ASM Listener Issue Affecting Flex ASM Functionality: Example 1-8
- Detecting and Resolving Memory Leaks: Example 1-9
- New crsctl Commands: Example 1-10
- CVU Enhancements: Diagnostics 1-12
- CVU Enhancements: Fixups 1-14
- CVU Enhancements: Error Reporting 1-15
- Autonomous Health Framework (AHF) Repository 1-17
- Review of CHM and CHA 1-18
- CHM Improved Diagnosability 1-20
- New oclumon Commands: Example 1-22
- FPP Overview: Background 1-24
- High-Level FPP Architecture 1-25
- FPP Archiving of Gold Images 1-26
- Summary 1-27

### 2 Oracle ASM

- Objectives 2-2
- On-Demand Scrubbing: Background 2-3
- Extent-Based Scrubbing 2-4
- Summary 2-5

### 3 Oracle RAC

- Objectives 3-2
- Simplified Database Deployment 3-4
- Automatic PDB Resource Management 3-5
- PDB Placement Choice: Fill or Spread 3-7
- Ranking and CPU 3-8
- Rank and CPU: Startup Order 3-9
- Failure Handling 3-10

|   |      |
|---|------|
| Last Node Standing  | 3-11 |
| Configuring a PDB with Cardinality and Rank: Example        | 3-12 |
| Oracle RAC Connection Management: Overview                  | 3-15 |
| Distribute Connections Based on Service Attribute           | 3-16 |
| Load Balancing Options in Oracle RAC: Example               | 3-18 |
| Performance Considerations                                  | 3-19 |
| Smart Connection Rebalance                                  | 3-20 |
| Testing Smart Connection Rebalance: Example                 | 3-21 |
| Outcome without Smart Connection Rebalance: Example         | 3-22 |
| Outcome with Smart Connection Rebalance: Example            | 3-23 |
| Index Block Contention: Background                          | 3-25 |
| Sequence Usage in Oracle RAC                                | 3-26 |
| Right-Growing Index Optimizations                           | 3-27 |
| Ordered Sequence Enhancement in Oracle RAC                  | 3-28 |
| Ordered Sequence Optimization in 19c: Example               | 3-29 |
| Ordered Sequence Optimization in 23ai: Example              | 3-30 |
| Review of Buddy Recovery for Reconfiguration                | 3-36 |
| Smooth Reconfiguration                                      | 3-37 |
| Single-Server Rolling Database Maintenance: Overview        | 3-39 |
| Local Rolling Database Maintenance: Overview                | 3-41 |
| Patching Oracle RAC Database in Local Rolling Mode: Example | 3-42 |
| Oracle RAC Two-Stage Rolling Patches: Overview              | 3-45 |
| Summary   | 3-46 |

#### **4 Oracle Data Guard Part 1**

|   |      |
|---|------|
| Objectives  | 4-2  |
| Oracle Active Data Guard                              | 4-3  |
| Automatic Primary Database Preparation                | 4-5  |
| Creating a Temp File after PDB Creation               | 4-7  |
| Automatic Temp File Creation on the Standby Database  | 4-8  |
| DML on Active Data Guard                              | 4-10 |
| Improved Performance of Redirected Transactions       | 4-11 |
| PL/SQL API for Data Guard Broker Management           | 4-13 |
| ORDS REST API for Data Guard Management               | 4-14 |
| Easier Checking of Data Guard Configurations          | 4-16 |
| Easier Checking of Fast-Start Failover Configurations | 4-17 |
| Easier Tracking of Role Transitions                   | 4-18 |
| Fast-Start Failover Lag Histogram                     | 4-19 |
| Enhanced Observer Diagnostic                          | 4-20 |
| Remote Snapshot Configuration for ADG Database        | 4-22 |
| Simplified AWR Snapshots on Active Data Guard         | 4-23 |

- Max Performance without Fast-Start Failover 4-25
- Max Performance with Fast-Start Failover 4-26
- Automatic Failover with MaxPerformance 4-31
- Minimized Stall in Fast-Start Failover Maximum Performance 4-33
- Choose the Lag Type for Maximum Performance Mode 4-35
- Strict Validation of Switchover Readiness 4-37
- New Command: VALIDATE DGConnectIdentifier 4-38
- New Commands: SHOW|EDIT ALL MEMBERS 4-39
- Summary 4-40

## **5 Oracle Data Guard Part 2**

- Objectives 5-2
- Traditional Per-CDB Data Guard Architecture 5-4
- Oracle Data Guard for Container Databases 5-5
- Data Guard Per PDB (DGPDB) 5-6
- DGPDB Integration Enhancements 5-7
- Current Limitations in DGPDB 5-8
- DGPDB Implementation: Overview 5-9
- Two Primary CDBs 5-10
- Prepare Both CDBs for Data Guard 5-11
- Set Up the Passwords in the Wallet 5-12
- Two Data Guard Configurations 5-13
- Add the Broker Configurations 5-14
- Internal Communication between CDBs 5-16
- Prepare the Configuration for DGPDB 5-17
- Add Target Pluggable Databases to Data Guard 5-18
- Check the PDBs Before Adding Them 5-19
- Add the PDBs to DGPDB 5-20
- Target Pluggable Databases Require Manual Copy 5-21
- Manual Copy: Example with scp on File System 5-22
- Manual Copy: Example with RMAN and ASM 5-23
- Transport Architecture 5-24
- Add the Standby Redo Logs to the Target PDBs 5-25
- The Apply Architecture 5-26
- Start the Redo Apply 5-27
- Show the Status of the PDBs 5-28
- PDB Switchover 5-29
- DGPDB and Transparent Data Encryption 5-30
- Copy the TDE Encryption Keys 5-31
- Application Continuity: Overview 5-33
- Use Cases 5-36

|                                      |      |
|--------------------------------------|------|
| What is Hybrid Data Guard?           | 5-41 |
| Hybrid DR Configurations: Challenges | 5-42 |
| Possible Solutions                   | 5-43 |
| How to Use It: Examples              | 5-44 |
| Summary                              | 5-45 |