

Oracle Grid Infrastructure 11g: Manage Clusterware and ASM

Student Guide - Volume I

D59999GC30

Edition 3.0

December 2012

D78226

ORACLE®

Authors

James Womack

James Spiller

Technical Contributors

David Brower

Jean-Francois Verrier

Mark Fuller

Mike Leatherman

Barb Lundhild

S. Matt Taylor

Rick Wessman

Joel Goodman

Harald Van Breederode

Markus Michalewicz

Technical Reviewers

Christopher Andrews

Christian Bauwens

Michael Cebulla

Jonathan Creighton

Al Flournoy

Andy Fortunak

Harald

Michael Hazel

Pete Jones

Jerry Lee

Markus Michalewicz

Peter Sharman

Ranbir Singh

Linda Smalley

Janet Stern

Richard Strohm

Branislav Valny

Doug Williams

Publishers

Veena Narasimhan

Srividya Rameshkumar

Jobi Varghese

Copyright © 2012, Oracle and/or its affiliates. All rights reserved.

Disclaimer

This document contains proprietary information and is protected by copyright and other intellectual property laws. You may copy and print this document solely for your own use in an Oracle training course. 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. If you find any problems in the document, please report them in writing to: Oracle University, 500 Oracle Parkway, Redwood Shores, California 94065 USA. This document is not warranted to be error-free.

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 RIGHTS

The U.S. Government's rights to use, modify, reproduce, release, perform, display, or disclose these training materials are restricted by the terms of the applicable Oracle license agreement and/or the applicable U.S. Government contract.

Trademark Notice

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

Contents

1 Oracle Grid Infrastructure Concepts

- Objectives 1-2
- Oracle Grid Infrastructure 1-3
- What Is a Cluster? 1-4
- What Is Clusterware? 1-5
- Oracle Clusterware 1-6
- Oracle Clusterware Architecture and Services 1-7
- Goals for Oracle Clusterware 1-8
- Oracle Clusterware Networking 1-9
- Grid Naming Service 1-11
- Single-Client Access Name 1-12
- Grid Plug and Play 1-14
- GPnP Domain 1-15
- GPnP Components 1-16
- GPnP Profile 1-17
- Oracle Automatic Storage Management (ASM) 1-18
- ASM: Key Features and Benefits 1-19
- ASM and Grid Infrastructure 1-20
- Quiz 1-21
- Summary 1-23

2 Oracle Clusterware Architecture

- Objectives 2-2
- Oracle Grid Infrastructure For a Cluster 2-3
- Oracle Cluster Registry (OCR) 2-4
- CSS Voting Disk Function 2-6
- Oracle Local Registry and High Availability 2-7
- Oracle Clusterware Initialization 2-8
- Clusterware Startup Details 2-9
- Clusterware Startup: The OHASD orarootagent 2-11
- Clusterware Startup Details: The CRSD orarootagent 2-13
- Clusterware Startup Details: The CRSD oraagent 2-14
- Clusterware Startup Details: The OHASD oraagent 2-15
- Controlling Oracle Clusterware 2-16
- Verifying the Status of Oracle Clusterware 2-17

Viewing the High Availability Services Stack 2-18
GPnP Architecture: Overview 2-19
How GPnP Works: Cluster Node Startup 2-21
How GPnP Works: Client Database Connections 2-22
Automatic Storage Management 2-23
Quiz 2-24
Summary 2-25

3 Grid Infrastructure Preinstallation Tasks

Objectives 3-2
Preinstallation Planning 3-3
Shared Storage Planning for Grid Infrastructure 3-4
Sizing Shared Storage for Oracle Clusterware 3-5
Storing the OCR in ASM 3-6
Managing Voting Disks in ASM 3-7
Installing ASMLib 3-8
Preparing ASMLib 3-9
Quiz 3-10
Grid Infrastructure Preinstallation Tasks 3-11
Oracle Grid Infrastructure 11g Installation 3-12
Checking System Requirements 3-13
Enabling the Name Service Cache Daemon (nscd) 3-14
Single-Client Access Name for the Cluster 3-15
Checking Network Requirements 3-16
IP Address Requirements with GNS 3-18
IP Address Requirements for Manual Configuration 3-19
Broadcast and Multicast Requirements 3-21
Interconnect NIC Guidelines 3-22
Redundant Interconnect Usage 3-23
Interconnect Link Aggregation: Single Switch 3-24
Interconnect Link Aggregation: Multiswitch 3-26
Additional Interconnect Guidelines 3-27
Software Requirements (Kernel) 3-28
32-Bit Software Requirements: Packages 3-29
64-Bit Software Requirements: Packages 3-30
Oracle Validated Configuration RPM 3-31
Creating Groups and Users 3-33
Creating Groups and Users: Example 3-34
Shell Settings for the Grid Infrastructure User 3-35
Quiz 3-37
Summary 3-38

4 Grid Infrastructure Installation

- Objectives 4-2
- Installing Grid Infrastructure 4-3
- Choosing an Installation Type 4-4
- Grid Plug and Play Support 4-5
- Cluster Node Information 4-6
- Specify Network Interface Usage 4-7
- Storage Option Information 4-8
- Specify Cluster Configuration: Typical Installation 4-9
- Install Locations: Typical Installation 4-11
- Failure Isolation Support with IPMI 4-12
- Privileged Operating System Groups 4-14
- Installation and Inventory Locations 4-15
- Prerequisite Checks 4-16
- Finishing the Installation 4-17
- Verifying the Grid Infrastructure Installation 4-19
- Modifying Oracle Clusterware Binaries After Installation 4-20
- Quiz 4-22
- Summary 4-23
- Practice 4: Overview 4-24

5 Adding and Removing Cluster Nodes

- Objectives 5-2
- Adding Oracle Clusterware Homes 5-3
- Prerequisite Steps for Running addNode.sh 5-4
- Adding a Node with addNode.sh 5-6
- Completing OUI Node Addition 5-7
- Removing a Node from the Cluster 5-8
- Deleting a Node from the Cluster 5-9
- Deleting a Node from a Cluster (GNS in Use) 5-12
- Deleting a Node from the Cluster 5-13
- Quiz 5-14
- Summary 5-15
- Practice 5: Overview 5-16

6 Administering Oracle Clusterware

- Objectives 6-2
- Managing Oracle Clusterware 6-3
- Managing Clusterware with Enterprise Manager 6-4
- Controlling Oracle Clusterware 6-5
- Controlling Oracle High Availability Services 6-6

- Verifying the Status of Oracle Clusterware 6-7
- Determining the Location of Oracle Clusterware Configuration Files 6-8
- Checking the Integrity of Oracle Clusterware Configuration Files 6-9
- Backing Up and Recovering the Voting Disk 6-10
- Adding, Deleting, or Migrating Voting Disks 6-11
- Locating the OCR Automatic Backups 6-12
- Changing the Automatic OCR Backup Location 6-13
- Adding, Replacing, and Repairing OCR Locations 6-14
- Removing an Oracle Cluster Registry Location 6-15
- Migrating OCR Locations to ASM 6-16
- Migrating OCR from ASM to Other Shared Storage 6-17
- Performing Manual OCR Backups 6-18
- Recovering the OCR by Using Physical Backups 6-19
- Oracle Local Registry 6-20
- Oracle Interface Configuration Tool: oifcfg 6-22
- Determining the Current Network Settings 6-23
- Configuring Redundant Interconnect Usage with oifcfg 6-24
- Changing the Public VIP Addresses for Non-GPnP Clusters 6-25
- Changing the Interconnect Adapter 6-27
- Managing SCAN VIP and SCAN Listener Resources 6-29
- Quiz 6-33
- Summary 6-35
- Practice 6: Overview 6-36

7 Upgrading and Patching Grid Infrastructure

- Objectives 7-2
- Out-of-Place Oracle Clusterware Upgrade 7-3
- Oracle Clusterware Upgrade 7-4
- Types of Patches 7-5
- Patch Properties 7-7
- Configuring the Software Library 7-8
- Setting Up Patching 7-9
- Starting the Provisioning Daemon 7-10
- Obtaining Oracle Clusterware Patches 7-11
- Downloading Patches 7-14
- Rolling Patches 7-15
- Checking Software Versions 7-16
- Installing a Rolling Patchset with OUI 7-17
- Patchset OUI 7-18
- Installing a Rolling Patchset with OUI 7-19
- OPatch: Overview 7-20

OPatch: General Usage 7-21
Before Patching with OPatch 7-22
Installing a Rolling Patch with OPatch 7-23
Quiz 7-28
Summary 7-30
Practice 7: Overview 7-31

8 Troubleshooting Oracle Clusterware

Objectives 8-2
Golden Rule in Debugging Oracle Clusterware 8-3
Monitoring Oracle Clusterware 8-5
Cluster Health Monitor (CHM) 8-7
oclumon Utility 8-8
oclumon debug Command 8-9
clumon dumpnodeview Command 8-10
oclumon dumpnodeview Command 8-11
oclumon manage Command 8-12
Oracle Clusterware Main Log Files 8-13
Oracle Clusterware Alerts 8-14
Diagnostic Record Unique IDs (DRUIDs) 8-15
Diagnostics Collection Script 8-16
Cluster Verify: Overview 8-17
Cluster Verify Components 8-18
Cluster Verify Locations 8-19
Cluster Verify Configuration File 8-20
Cluster Verify Output: Example 8-22
Enabling Resource Debugging 8-23
Dynamic Debugging 8-25
Enabling Tracing for Java-Based Tools 8-27
Preserving Log Files Before Wrapping 8-28
Process Roles for Node Reboots 8-29
Determining Which Process Caused Reboot 8-30
Using ocrdump to View Logical Contents of the OCR 8-31
Checking the Integrity of the OCR 8-32
OCR-Related Tools for Debugging 8-33
Browsing My Oracle Support Knowledge Articles 8-35
Quiz 8-36
Summary 8-37
Practice 8: Overview 8-38

9 Making Applications Highly Available with Oracle Clusterware

- Objectives 9-2
- Oracle Clusterware High Availability (HA) 9-3
- Oracle Clusterware HA Components 9-4
- Resource Management Options 9-5
- Server Pools 9-6
- Server Pool Attributes 9-7
- GENERIC and FREE Server Pools 9-9
- Assignment of Servers to Server Pools 9-11
- Server Attributes and States 9-12
- Creating Server Pools with srvctl and crsctl 9-14
- Managing Server Pools with srvctl and crsctl 9-15
- Adding Server Pools with Enterprise Manager 9-16
- Managing Server Pools with Enterprise Manager 9-17
- Clusterware Resource Modeling 9-18
- Resource Types 9-20
- Adding a Resource Type 9-21
- Resource Type Parameters 9-23
- Resource Type Advanced Settings 9-24
- Defining Resource Dependencies 9-25
- Creating an Application VIP by Using crsctl 9-27
- Creating an Application VIP by Using EM 9-29
- Managing Clusterware Resources with EM 9-30
- Adding Resources with EM 9-31
- Adding Resources by Using crsctl 9-36
- Managing Resources with EM 9-37
- Managing Resources with crsctl 9-40
- HA Events: ONS and FAN 9-42
- Managing Oracle Notification Server with srvctl 9-43
- Quiz 9-44
- Summary 9-46
- Practice 9: Overview 9-47

10 ASM: Overview

- Objectives 10-2
- What Is Oracle ASM? 10-3
- ASM and ASM Cluster File System 10-4
- ASM Key Features and Benefits 10-6
- ASM Instance Designs: Nonclustered ASM and Oracle Databases 10-7
- ASM Instance Designs: Clustered ASM for Clustered Databases 10-8
- ASM Instance Designs: Clustered ASM for Mixed Databases 10-9

- ASM System Privileges 10-10
- ASM OS Groups with Role Separation 10-11
- Authentication for Accessing ASM Instances 10-12
- Password-Based Authentication for ASM 10-13
- Managing the ASM Password File 10-14
- Using a Single OS Group 10-15
- Using Separate OS Groups 10-16
- ASM Components: Software 10-17
- ASM Components: ASM Instance 10-18
- ASM Components: ASM Instance Primary Processes 10-20
- ASM Components: Node Listener 10-21
- ASM Components: Configuration Files 10-22
- ASM Components: Group Services 10-23
- ASM Components: ASM Disk Group 10-24
- ASM Disk Group: Failure Groups 10-25
- ASM Components: ASM Disks 10-26
- ASM Components: ASM Files 10-27
- ASM Files: Extents and Striping 10-28
- ASM Files: Mirroring 10-29
- ASM Components: ASM Clients 10-30
- ASM Components: ASM Utilities 10-31
- ASM Scalability 10-33
- Quiz 10-35
- Summary 10-36

11 Administering ASM

- Objectives 11-2
- Managing ASM with ASMCA 11-3
- Starting and Stopping ASM Instances by Using ASMCA and ASMCMD 11-4
- Starting and Stopping ASM Instances by Using srvctl 11-5
- Starting and Stopping ASM Instances by Using SQL*Plus 11-6
- Starting and Stopping ASM Instances Containing Cluster Files 11-8
- ASM Initialization Parameters 11-9
- ASM_DISKGROUPS 11-10
- Disk Groups Mounted at Startup 11-11
- ASM_DISKSTRING 11-12
- ASM_POWER_LIMIT 11-14
- INSTANCE_TYPE 11-15
- CLUSTER_DATABASE 11-16
- MEMORY_TARGET 11-17
- Adjusting ASM Instance Parameters in SPFILEs 11-18

Starting and Stopping the Node Listener 11-19
ASM Dynamic Performance Views 11-20
ASM Dynamic Performance Views Diagram 11-21
Quiz 11-23
Summary 11-25
Practice 11 Overview: Administering ASM Instances 11-26

12 Administering ASM Disk Groups

Objectives 12-2
Disk Group: Overview 12-3
Creating a New Disk Group 12-4
Creating a New Disk Group with ASMCMD 12-6
Creating an ASM Disk Group with ASMCA 12-7
Creating an ASM Disk Group: Advanced Options 12-8
Creating a Disk Group with Enterprise Manager 12-9
Disk Group Attributes 12-11
V\$ASM_ATTRIBUTE 12-13
Compatibility Attributes 12-14
Features Enabled by Disk Group Compatibility Attributes 12-15
Support for 4 KB Sector Disk Drives 12-16
Supporting 4 KB Sector Disks 12-17
ASM Support for 4 KB Sector Disks 12-18
Using the SECTOR_SIZE Clause 12-19
Viewing ASM Disk Groups 12-21
Viewing ASM Disk Information 12-23
Extending an Existing Disk Group 12-25
Dropping Disks from an Existing Disk Group 12-26
REBALANCE POWER 0 12-27
V\$ASM_OPERATION 12-28
Adding and Dropping in the Same Command 12-30
Adding and Dropping Failure Groups 12-31
Undropping Disks in Disk Groups 12-32
Mounting and Dismounting Disk Groups 12-33
Viewing Connected Clients 12-34
Dropping Disk Groups 12-35
Checking the Consistency of Disk Group Metadata 12-36
ASM Fast Mirror Resync 12-37
Preferred Read Failure Groups 12-38
Preferred Read Failure Groups: Best Practice 12-39
Viewing ASM Disk Statistics 12-40
Performance, Scalability, and Manageability Considerations for Disk Groups 12-42

Quiz 12-43
Summary 12-45
Practice 12 Overview: Administering ASM Disk Groups 12-46

13 Administering ASM Files, Directories, and Templates

Objectives 13-2
ASM Clients 13-3
Interaction Between Database Instances and ASM 13-5
Accessing ASM Files by Using RMAN 13-6
Accessing ASM Files by Using XML DB 13-8
Accessing ASM Files by Using DBMS_FILE_TRANSFER 13-9
Accessing ASM Files by Using ASMCMD 13-10
Fully Qualified ASM File Names 13-11
Other ASM File Names 13-13
Valid Contexts for the ASM File Name Forms 13-15
Single File Creation: Examples 13-16
Multiple File Creation: Example 13-17
View ASM Aliases, Files, and Directories 13-18
Viewing ASM Files 13-20
ASM Directories 13-21
Managing ASM Directories 13-22
Managing Alias File Names 13-23
Disk Group Templates 13-24
Viewing Templates 13-26
Managing Disk Group Templates 13-27
Managing Disk Group Templates with ASMCMD 13-28
Using Disk Group Templates 13-29
ASM Intelligent Data Placement 13-30
Guidelines for Intelligent Data Placement 13-31
Assigning Files to Disk Regions 13-32
Assigning Files to Disk Regions with Enterprise Manager 13-33
Monitoring Intelligent Data Placement 13-34
ASM Access Control Lists 13-35
ASM ACL Prerequisites 13-36
Managing ASM ACL with SQL Commands 13-37
Managing ASM ACL with ASMCMD Commands 13-38
Managing ASM ACL with Enterprise Manager 13-39
ASM ACL Guidelines 13-41
Quiz 13-42
Summary 13-44
Practice 13: Overview 13-45

14 Administering ASM Cluster File Systems

- Objectives 14-2
- ASM Files and Volumes 14-3
- ACFS and ADVM Architecture: Overview 14-4
- ADVM Processes 14-6
- ADVM Restrictions 14-7
- ASM Cluster File System 14-8
- ADVM Space Allocation 14-9
- Striping Inside the Volume 14-10
- Volume Striping: Example 14-11
- Creating an ACFS Volume 14-13
- Creating an ASM Dynamic Volume with Enterprise Manager 14-14
- Managing ADVM Dynamic Volumes 14-17
- Creating an ASM Cluster File System with Enterprise Manager 14-18
- Managing Dynamic Volumes with SQL*PLUS 14-19
- Registering an ACFS Volume 14-20
- Creating an ACFS Volume with ASMCA 14-21
- Creating the ACFS File System with ASMCA 14-22
- Mounting the ACFS File System with ASMCA 14-23
- Managing ACFS with EM 14-24
- Extending ASMCMDB for Dynamic Volumes 14-25
- Linux-UNIX File System APIs 14-26
- Linux-UNIX Extensions 14-27
- ACFS Platform-Independent Commands 14-28
- ACFS Snapshots 14-29
- Managing ACFS Snapshots 14-30
- Managing ACFS Snapshots with Enterprise Manager 14-32
- Creating ACFS Snapshots 14-33
- Managing Snapshots 14-34
- Viewing Snapshots 14-35
- ACFS Replication 14-36
- ACFS Replication Requirements 14-38
- Managing ACFS Replication 14-40
- ACFS Backups 14-43
- ACFS Performance 14-44
- Using ACFS Volumes After Reboot 14-45
- ACFS Views 14-46
- Quiz 14-47
- Summary 14-48
- Practice 14: Overview 14-49

A Practices and Solutions

B DHCP and DNS Configuration for GNS

- Objectives B-2
- GNS: Overview B-3
- DHCP Service B-4
- DHCP Configuration: Example B-5
- DHCP Configuration Example B-6
- DNS Concepts B-7
- DNS Forwarding for GNS B-9
- DNS Configuration: Example B-11
- DNS Configuration: Detail B-13

C Cloning Grid Infrastructure

- Objectives C-2
- What Is Cloning? C-3
- Benefits of Cloning Grid Infrastructure C-4
- Creating a Cluster by Cloning Grid Infrastructure C-5
- Preparing the Oracle Clusterware Home for Cloning C-6
- Cloning to Create a New Oracle Clusterware Environment C-9
- clone.pl Script C-13
- clone.pl Environment Variables C-14
- clone.pl Command Options C-15
- Cloning to Create a New Oracle Clusterware Environment C-16
- Log Files Generated During Cloning C-18
- Cloning to Extend Oracle Clusterware to More Nodes C-19
- Quiz C-25
- Summary C-27

D RAC Concepts

- Objectives D-2
- Benefits of Using RAC D-3
- Clusters and Scalability D-4
- Levels of Scalability D-5
- Scaleup and Speedup D-6
- Speedup/Scaleup and Workloads D-7
- I/O Throughput Balanced: Example D-8
- Performance of Typical Components D-9
- Necessity of Global Resources D-10
- Global Resources Coordination D-11
- Global Cache Coordination: Example D-12

Write to Disk Coordination: Example	D-13
Dynamic Reconfiguration	D-14
Object Affinity and Dynamic Remastering	D-15
Global Dynamic Performance Views	D-16
Efficient Internode Row-Level Locking	D-17
Parallel Execution with RAC	D-18
Summary	D-19