

Oracle® Grid Infrastructure

Installation and Upgrade Guide

12c Release 2 (12.2) for Linux

E49640-11

July 2017

Oracle Grid Infrastructure Installation and Upgrade Guide, 12c Release 2 (12.2) for Linux

E49640-11

Copyright © 2013, 2017, Oracle and/or its affiliates. All rights reserved.

Primary Author: Aparna Kamath

Contributing Authors: Douglas Williams, Mark Bauer, Janet Stern

Contributors: Prakash Jashnani, Markus Michalewicz, Aneesh Khandelwal, James Williams, Ian Cookson, Jonathan Creighton, Angad Gokakkar, Srinivas Poovala, Khethavath Singh, Prasad Bagal, Balaji Pagadala, Neha Avasthy, Mark Scardina, Hanlin Chien, Anil Nair, Apparsamy Perumal, Samarjeet Tomar, Akshay Shah, Kevin Jernigan, Eric Belden, Mark Fuller, Barbara Glover, Saar Maoz, Binoy Sukumaran

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein 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.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and 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.

Contents

Preface	xvii
Audience	xvii
Documentation Accessibility	xvii
Related Documentation	xviii
Conventions.....	xviii
Changes in This Release for Oracle Grid Infrastructure.....	xix
Changes in Oracle Grid Infrastructure 12c Release 2 (12.2)	xix
New Features	xix
Deprecated Features.....	xxvi
Desupported Features	xxvi
Changes in Oracle Grid Infrastructure 12c Release 1 (12.1)	xxvii
New Features for Oracle Grid Infrastructure 12c Release 1 (12.1.0.2)	xxvii
New Features for Oracle Grid Infrastructure 12c Release 1 (12.1.0.1)	xxviii
Deprecated Features.....	xxxi
Desupported Features	xxxi
1 Oracle Grid Infrastructure Installation Checklist	
1.1 Server Hardware Checklist for Oracle Grid Infrastructure.....	1-1
1.2 Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC.....	1-2
1.3 Server Configuration Checklist for Oracle Grid Infrastructure	1-4
1.4 Network Checklist for Oracle Grid Infrastructure	1-5
1.5 User Environment Configuration Checklist for Oracle Grid Infrastructure.....	1-7
1.6 Storage Checklist for Oracle Grid Infrastructure	1-9
1.7 Cluster Deployment Checklist for Oracle Grid Infrastructure	1-11
1.8 Installer Planning Checklist for Oracle Grid Infrastructure.....	1-12
2 Checking and Configuring Server Hardware for Oracle Grid Infrastructure	
2.1 Logging In to a Remote System Using X Window System.....	2-1
2.2 Checking Server Hardware and Memory Configuration.....	2-2

3	Automatically Configuring Oracle Linux with Oracle Preinstallation RPM	
3.1	Overview of Oracle Linux Configuration with Oracle RPMs	3-1
3.2	Installing the Oracle Preinstallation RPM with ULN Support	3-2
3.3	Installing the Oracle Preinstallation RPM From Unbreakable Linux Network	3-3
3.4	Installing a New Oracle Linux Installation from DVDs or Images	3-4
3.5	Installing Oracle Linux with Oracle Linux Yum Server Support	3-5
3.6	Configuring Oracle Ksplice to Perform Oracle Linux Updates	3-6
3.7	Configure Additional Operating System Features	3-6
4	Configuring Operating Systems for Oracle Grid Infrastructure on Linux	
4.1	Guidelines for Linux Operating System Installation.....	4-2
4.1.1	Completing a Minimal Linux Installation	4-3
4.1.2	Completing a Default Linux Installation	4-4
4.1.3	About Oracle Linux with the Unbreakable Enterprise Kernel	4-4
4.1.4	About the Oracle Preinstallation RPM.....	4-5
4.1.5	Restrictions for HugePages and Transparent HugePages Configurations.....	4-6
4.2	Reviewing Operating System and Software Upgrade Best Practices	4-7
4.2.1	General Upgrade Best Practices	4-7
4.2.2	New Server Operating System Upgrade Option.....	4-8
4.2.3	Oracle ASM Upgrade Notifications.....	4-8
4.3	Reviewing Operating System Security Common Practices.....	4-9
4.4	About Installation Fixup Scripts.....	4-9
4.5	About Operating System Requirements.....	4-10
4.6	Using Oracle RPM Checker on IBM: Linux on System z.....	4-10
4.7	Operating System Requirements for x86-64 Linux Platforms	4-11
4.7.1	Supported Oracle Linux 7 Distributions for x86-64	4-12
4.7.2	Supported Oracle Linux 6 Distributions for x86-64	4-13
4.7.3	Supported Red Hat Enterprise Linux 7 Distributions for x86-64.....	4-15
4.7.4	Supported Red Hat Enterprise Linux 6 Distributions for x86-64.....	4-16
4.7.5	Supported SUSE Linux Enterprise Server 12 Distributions for x86-64.....	4-17
4.8	Operating System Requirements for IBM: Linux on System z	4-20
4.8.1	Supported Red Hat Enterprise Linux 7 Distributions for IBM: Linux on System z..	4-20
4.8.2	Supported Red Hat Enterprise Linux 6 Distributions for IBM: Linux on System z..	4-21
4.8.3	Supported SUSE Linux Enterprise Server 12 Distributions for IBM: Linux on System z.....	4-22
4.9	Additional Drivers and Software Packages for Linux	4-23
4.9.1	Installing PAM for Login Authentication on Linux.....	4-24
4.9.2	Installation Requirements for OCFS2.....	4-24
4.9.3	Installation Requirements for ODBC and LDAP.....	4-25
4.9.4	Installing Oracle Messaging Gateway.....	4-26
4.9.5	Installation Requirements for Programming Environments for Linux.....	4-26
4.9.6	Installation Requirements for Web Browsers.....	4-27

4.10	Checking Kernel and Package Requirements for Linux	4-27
4.11	Installing the cvuqdisk RPM for Linux	4-28
4.12	Reviewing HugePages Memory Allocation.....	4-28
4.13	Disabling Transparent HugePages.....	4-29
4.14	Enabling the Name Service Cache Daemon	4-30
4.15	Verifying the Disk I/O Scheduler on Linux	4-30
4.16	Using Automatic SSH Configuration During Installation.....	4-31
4.17	Setting Network Time Protocol for Cluster Time Synchronization	4-32

5 Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

5.1	About Oracle Grid Infrastructure Network Configuration Options.....	5-2
5.2	Understanding Network Addresses	5-3
5.2.1	About the Public IP Address	5-3
5.2.2	About the Private IP Address.....	5-3
5.2.3	About the Virtual IP Address.....	5-4
5.2.4	About the Grid Naming Service (GNS) Virtual IP Address	5-4
5.2.5	About the SCAN.....	5-5
5.3	Network Interface Hardware Minimum Requirements	5-6
5.4	Private IP Interface Configuration Requirements.....	5-7
5.5	IPv4 and IPv6 Protocol Requirements	5-8
5.6	Oracle Grid Infrastructure IP Name and Address Requirements	5-9
5.6.1	About Oracle Grid Infrastructure Name Resolution Options	5-10
5.6.2	Cluster Name and SCAN Requirements	5-11
5.6.3	IP Name and Address Requirements For Grid Naming Service (GNS)	5-11
5.6.4	IP Name and Address Requirements For Multi-Cluster GNS	5-11
5.6.5	IP Name and Address Requirements for Manual Configuration of Cluster.....	5-13
5.6.6	Confirming the DNS Configuration for SCAN.....	5-15
5.7	Broadcast Requirements for Networks Used by Oracle Grid Infrastructure.....	5-15
5.8	Multicast Requirements for Networks Used by Oracle Grid Infrastructure	5-16
5.9	Domain Delegation to Grid Naming Service.....	5-16
5.9.1	Choosing a Subdomain Name for Use with Grid Naming Service	5-16
5.9.2	Configuring DNS for Cluster Domain Delegation to Grid Naming Service	5-17
5.10	Configuration Requirements for Oracle Flex Clusters.....	5-18
5.10.1	Understanding Oracle Flex Clusters	5-18
5.10.2	About Oracle Flex ASM Clusters Networks.....	5-19
5.10.3	General Requirements for Oracle Flex Cluster Configuration	5-21
5.10.4	Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses	5-21
5.10.5	Oracle Flex Cluster Manually-Assigned Addresses.....	5-21
5.11	Grid Naming Service Cluster Configuration Example	5-22
5.12	Manual IP Address Configuration Example	5-24
5.13	Network Interface Configuration Options.....	5-25
5.14	Multiple Private Interconnects and Oracle Linux.....	5-25

6 Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle Database

6.1	Creating Groups, Users and Paths for Oracle Grid Infrastructure.....	6-1
6.1.1	Determining If an Oracle Inventory and Oracle Inventory Group Exist	6-2
6.1.2	Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist.....	6-3
6.1.3	About Oracle Installation Owner Accounts	6-4
6.1.4	Restrictions for Oracle Software Installation Owners.....	6-4
6.1.5	Identifying an Oracle Software Owner User Account	6-5
6.1.6	About the Oracle Base Directory for the Grid User	6-6
6.1.7	About the Oracle Home Directory for Oracle Grid Infrastructure Software.....	6-6
6.1.8	About Creating the Oracle Home and Oracle Base Directory	6-7
6.2	Oracle Installations with Standard and Job Role Separation Groups and Users	6-8
6.2.1	About Oracle Installations with Job Role Separation.....	6-8
6.2.2	Standard Oracle Database Groups for Database Administrators	6-9
6.2.3	Extended Oracle Database Groups for Job Role Separation	6-9
6.2.4	Creating an ASMSNMP User	6-10
6.2.5	Oracle Automatic Storage Management Groups for Job Role Separation.....	6-11
6.3	Creating Operating System Privileges Groups	6-11
6.3.1	Creating the OSASM Group	6-12
6.3.2	Creating the OSDBA for ASM Group	6-12
6.3.3	Creating the OSOPER for ASM Group	6-13
6.3.4	Creating the OSDBA Group for Database Installations	6-13
6.3.5	Creating an OSOPER Group for Database Installations.....	6-13
6.3.6	Creating the OSBACKUPDBA Group for Database Installations.....	6-14
6.3.7	Creating the OSDGDBA Group for Database Installations	6-14
6.3.8	Creating the OSKMDBA Group for Database Installations.....	6-14
6.3.9	Creating the OSRACDBA Group for Database Installations.....	6-14
6.4	Creating Operating System Oracle Installation User Accounts.....	6-14
6.4.1	Creating an Oracle Software Owner User	6-15
6.4.2	Modifying Oracle Owner User Groups.....	6-15
6.4.3	Identifying Existing User and Group IDs.....	6-16
6.4.4	Creating Identical Database Users and Groups on Other Cluster Nodes.....	6-16
6.4.5	Example of Creating Minimal Groups, Users, and Paths.....	6-17
6.4.6	Example of Creating Role-allocated Groups, Users, and Paths	6-19
6.5	Configuring Grid Infrastructure Software Owner User Environments	6-22
6.5.1	Environment Requirements for Oracle Software Owners	6-22
6.5.2	Procedure for Configuring Oracle Software Owner Environments	6-23
6.5.3	Checking Resource Limits for Oracle Software Installation Users	6-25
6.5.4	Setting Remote Display and X11 Forwarding Configuration.....	6-26
6.5.5	Preventing Installation Errors Caused by Terminal Output Commands	6-27
6.6	Enabling Intelligent Platform Management Interface (IPMI)	6-28
6.6.1	Requirements for Enabling IPMI	6-28

6.6.2	Configuring the IPMI Management Network	6-29
6.6.3	Configuring the Open IPMI Driver	6-29
6.6.4	Configuring the BMC.....	6-30
6.6.5	Example of BMC Configuration Using IPMItool.....	6-31
6.7	Determining Root Script Execution Plan.....	6-33
7	Supported Storage Options for Oracle Database and Oracle Grid Infrastructure	
7.1	Supported Storage Options for Oracle Grid Infrastructure.....	7-2
7.2	Oracle ACFS and Oracle ADVM	7-4
7.2.1	Oracle ACFS and Oracle ADVM Support on Linux x86-64	7-4
7.2.2	Restrictions and Guidelines for Oracle ACFS	7-5
7.3	Storage Considerations for Oracle Grid Infrastructure and Oracle RAC.....	7-6
7.4	Guidelines for Using Oracle ASM Disk Groups for Storage	7-7
7.5	Guidelines for Configuring Oracle ASM Disk Groups on NFS.....	7-8
7.6	Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC.....	7-10
7.7	Using a Cluster File System for Oracle Clusterware Files	7-10
7.8	About NFS Storage for Data Files	7-10
7.9	About Direct NFS Client Mounts to NFS Storage Devices.....	7-11
8	Configuring Storage for Oracle Grid Infrastructure	
8.1	Configuring Storage for Oracle Automatic Storage Management	8-2
8.1.1	Identifying Storage Requirements for Oracle Automatic Storage Management.....	8-3
8.1.2	Oracle Clusterware Storage Space Requirements	8-7
8.1.3	About the Grid Infrastructure Management Repository.....	8-10
8.1.4	Using an Existing Oracle ASM Disk Group	8-10
8.1.5	About Upgrading Existing Oracle Automatic Storage Management Instances	8-11
8.1.6	Selecting Disks to use with Oracle ASM Disk Groups	8-11
8.1.7	Specifying the Oracle ASM Disk Discovery String	8-12
8.1.8	Creating Files on a NAS Device for Use with Oracle Automatic Storage Management.....	8-12
8.2	Configuring Storage Device Path Persistence Using Oracle ASMFD.....	8-13
8.2.1	About Oracle ASM with Oracle ASM Filter Driver	8-14
8.3	Using Disk Groups with Oracle Database Files on Oracle ASM.....	8-14
8.3.1	Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM.....	8-15
8.3.2	Creating Disk Groups for Oracle Database Data Files.....	8-15
8.3.3	Creating Directories for Oracle Database Files	8-15
8.4	Configuring File System Storage for Oracle Database	8-16
8.4.1	Configuring NFS Buffer Size Parameters for Oracle Database	8-17
8.4.2	Checking TCP Network Protocol Buffer for Direct NFS Client	8-17
8.4.3	Creating an oranfstab File for Direct NFS Client.....	8-18
8.4.4	Enabling and Disabling Direct NFS Client Control of NFS	8-20
8.4.5	Enabling Hybrid Columnar Compression on Direct NFS Client.....	8-21
8.5	Creating Member Cluster Manifest File for Oracle Member Clusters.....	8-21

8.6	Configuring Oracle Automatic Storage Management Cluster File System	8-22
8.7	Checking OCFS2 Version Manually	8-23

9 Installing Oracle Grid Infrastructure

9.1	About Image-Based Oracle Grid Infrastructure Installation.....	9-2
9.2	Understanding Cluster Configuration Options	9-3
9.2.1	About Oracle Standalone Clusters.....	9-3
9.2.2	About Oracle Cluster Domain and Oracle Domain Services Cluster.....	9-3
9.2.3	About Oracle Member Clusters.....	9-5
9.2.4	About Oracle Extended Clusters.....	9-6
9.3	Installing Oracle Grid Infrastructure for a New Cluster.....	9-7
9.3.1	Installing Oracle Standalone Cluster.....	9-7
9.3.2	Installing Oracle Domain Services Cluster	9-14
9.3.3	Installing Oracle Member Clusters	9-21
9.4	Installing Oracle Grid Infrastructure Using a Cluster Configuration File	9-26
9.5	Installing Only the Oracle Grid Infrastructure Software	9-28
9.5.1	Installing Software Binaries for Oracle Grid Infrastructure for a Cluster.....	9-28
9.5.2	Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster.....	9-29
9.5.3	Configuring the Software Binaries Using a Response File.....	9-30
9.5.4	Setting Ping Targets for Network Checks.....	9-30
9.6	About Deploying Oracle Grid Infrastructure Using Rapid Home Provisioning	9-31
9.7	Confirming Oracle Clusterware Function.....	9-32
9.8	Confirming Oracle ASM Function for Oracle Clusterware Files.....	9-32
9.9	Understanding Offline Processes in Oracle Grid Infrastructure	9-33

10 Oracle Grid Infrastructure Postinstallation Tasks

10.1	Required Postinstallation Tasks	10-1
10.1.1	Downloading and Installing Patch Updates	10-2
10.2	Recommended Postinstallation Tasks	10-2
10.2.1	Tuning Semaphore Parameters	10-3
10.2.2	Creating a Backup of the root.sh Script.....	10-3
10.2.3	Downloading and Installing the ORAchk Health Check Tool.....	10-3
10.2.4	Creating a Fast Recovery Area	10-4
10.2.5	Checking the SCAN Configuration	10-5
10.2.6	Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications.....	10-6
10.3	About Changes in Default SGA Permissions for Oracle Database	10-7
10.4	Using Earlier Oracle Database Releases with Oracle Grid Infrastructure.....	10-7
10.4.1	General Restrictions for Using Earlier Oracle Database Releases.....	10-8
10.4.2	Configuring Earlier Release Oracle Database on Oracle ACFS.....	10-8
10.4.3	Managing Server Pools with Earlier Database Versions	10-9
10.4.4	Making Oracle ASM Available to Earlier Oracle Database Releases	10-10
10.4.5	Using ASMCA to Administer Disk Groups for Earlier Database Releases.....	10-10

10.4.6	Using the Correct LSNRCTL Commands.....	10-10
10.5	Modifying Oracle Clusterware Binaries After Installation.....	10-10

11 Upgrading Oracle Grid Infrastructure

11.1	Understanding Out-of-Place Upgrade	11-2
11.2	About Oracle Grid Infrastructure Upgrade and Downgrade	11-3
11.3	Options for Oracle Grid Infrastructure Upgrades	11-3
11.4	Restrictions for Oracle Grid Infrastructure Upgrades.....	11-4
11.5	Preparing to Upgrade an Existing Oracle Clusterware Installation.....	11-6
11.5.1	Upgrade Checklist for Oracle Grid Infrastructure	11-6
11.5.2	Checks to Complete Before Upgrading Oracle Grid Infrastructure	11-8
11.5.3	Moving Oracle Clusterware Files from NFS to Oracle ASM	11-9
11.5.4	Running the Oracle ORAchk Upgrade Readiness Assessment.....	11-10
11.5.5	Using CVU to Validate Readiness for Oracle Clusterware Upgrades.....	11-11
11.6	Understanding Rolling Upgrades Using Batches	11-12
11.7	Performing Rolling Upgrade of Oracle Grid Infrastructure.....	11-13
11.7.1	Upgrading Oracle Grid Infrastructure from an Earlier Release	11-13
11.7.2	Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable	11-16
11.7.3	Joining Inaccessible Nodes After Forcing an Upgrade.....	11-16
11.7.4	Changing the First Node for Install and Upgrade.....	11-16
11.8	About Upgrading Oracle Grid Infrastructure Using Rapid Home Provisioning	11-17
11.9	Applying Patches to Oracle Grid Infrastructure	11-18
11.9.1	About Individual (One-Off) Oracle Grid Infrastructure Patches	11-18
11.9.2	About Oracle Grid Infrastructure Software Patch Levels	11-18
11.9.3	Patching Oracle Grid Infrastructure to a Software Patch Level	11-19
11.10	Updating Oracle Enterprise Manager Cloud Control Target Parameters.....	11-19
11.10.1	Updating the Enterprise Manager Cloud Control Target After Upgrades	11-20
11.10.2	Updating the Enterprise Manager Agent Base Directory After Upgrades.....	11-20
11.10.3	Registering Resources with Oracle Enterprise Manager After Upgrades	11-20
11.11	Unlocking the Existing Oracle Clusterware Installation.....	11-21
11.12	Checking Cluster Health Monitor Repository Size After Upgrading.....	11-22
11.13	Downgrading Oracle Clusterware After an Upgrade.....	11-23
11.13.1	Options for Oracle Grid Infrastructure Downgrades	11-23
11.13.2	Restrictions for Oracle Grid Infrastructure Downgrades.....	11-24
11.13.3	Downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1)	11-24
11.13.4	Downgrading to Oracle Grid Infrastructure 11g Release 2 (11.2).....	11-26
11.13.5	Downgrading Oracle Grid Infrastructure after Upgrade Fails.....	11-27
11.13.6	Downgrading Oracle Grid Infrastructure after Upgrade Fails on Remote Nodes	11-28
11.14	Completing Failed or Interrupted Installations and Upgrades	11-29
11.14.1	Completing Failed Installations and Upgrades	11-29
11.14.2	Continuing Incomplete Upgrade of First Nodes	11-30
11.14.3	Continuing Incomplete Upgrades on Remote Nodes.....	11-30
11.14.4	Continuing Incomplete Installation on First Node	11-31

11.14.5 Continuing Incomplete Installation on Remote Nodes	11-31
11.15 Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure ..	11-32

12 Removing Oracle Database Software

12.1 About Oracle Deinstallation Options	12-2
12.2 Oracle Deinstallation Tool (Deinstall)	12-4
12.3 Deinstallation Examples for Oracle Database	12-6
12.4 Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster	12-7
12.5 Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster	12-10
12.6 Relinking Oracle Grid Infrastructure for a Cluster Binaries	12-12
12.7 Changing the Oracle Grid Infrastructure Home Path	12-12
12.8 Unconfiguring Oracle Clusterware Without Removing Binaries.....	12-13
12.9 Unconfiguring Oracle Member Cluster.....	12-14

A Installing and Configuring Oracle Database Using Response Files

A.1 How Response Files Work	A-2
A.2 Reasons for Using Silent Mode or Response File Mode	A-2
A.3 Using Response Files.....	A-3
A.4 Preparing Response Files	A-3
A.4.1 Editing a Response File Template.....	A-4
A.4.2 Recording Response Files	A-5
A.5 Running Oracle Universal Installer Using a Response File.....	A-6
A.6 Running Configuration Assistants Using Response Files	A-7
A.6.1 Running Database Configuration Assistant Using Response Files	A-8
A.6.2 Running Net Configuration Assistant Using Response Files	A-9
A.7 Postinstallation Configuration Using Response File Created During Installation	A-10
A.7.1 Using the Installation Response File for Postinstallation Configuration.....	A-10
A.7.2 Running Postinstallation Configuration Using Response File	A-11
A.8 Postinstallation Configuration Using the ConfigToolAllCommands Script	A-12
A.8.1 About the Postinstallation Configuration File	A-13
A.8.2 Creating a Password Response File.....	A-14
A.8.3 Running Postinstallation Configuration Using a Password Response File.....	A-14

B Completing Preinstallation Tasks Manually

B.1 Configuring SSH Manually on All Cluster Nodes	B-1
B.1.1 Checking Existing SSH Configuration on the System.....	B-2
B.1.2 Configuring SSH on Cluster Nodes	B-2
B.1.3 Enabling SSH User Equivalency on Cluster Nodes.....	B-4
B.2 Configuring Storage Device Path Persistence Using Oracle ASMLIB.....	B-5
B.2.1 About Oracle ASM with Oracle ASMLIB.....	B-5
B.2.2 Installing and Configuring Oracle ASMLIB Software.....	B-6
B.2.3 Configuring Disk Devices to Use Oracle ASMLIB.....	B-8
B.2.4 Administering Oracle ASMLIB and Disks	B-10

B.2.5	Configuring Oracle ASMLIB for Multipath Disks	B-11
B.2.6	Deinstalling Oracle ASMLIB On Oracle Grid Infrastructure	B-14
B.3	Configuring Storage Device Path Persistence Manually	B-14
B.3.1	Configuring Device Persistence Manually for Oracle ASM	B-15
B.4	Configuring Kernel Parameters for Linux	B-17
B.4.1	Minimum Parameter Settings for Installation	B-17
B.4.2	Changing Kernel Parameter Values	B-19
B.4.3	Configuring Additional Kernel Settings for SUSE Linux	B-20
B.4.4	Setting UDP and TCP Kernel Parameters Manually	B-21

C Optimal Flexible Architecture

C.1	About the Optimal Flexible Architecture Standard	C-1
C.2	About Multiple Oracle Homes Support	C-2
C.3	About the Oracle Inventory Directory and Installation	C-3
C.4	Oracle Base Directory Naming Convention	C-4
C.5	Oracle Home Directory Naming Convention	C-4
C.6	Optimal Flexible Architecture File Path Examples	C-5

Index

List of Figures

9-1	Oracle Cluster Domain.....	9-4
-----	----------------------------	-----

List of Tables

1-1	Server Hardware Checklist for Oracle Grid Infrastructure.....	1-2
1-2	Operating System General Checklist for Oracle Grid Infrastructure and Oracle RAC....	1-3
1-3	Server Configuration Checklist for Oracle Grid Infrastructure.....	1-4
1-4	Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC.....	1-5
1-5	User Environment Configuration for Oracle Grid Infrastructure.....	1-8
1-6	Oracle Grid Infrastructure Storage Configuration Checks.....	1-9
1-7	Oracle Grid Infrastructure Cluster Deployment Checklist.....	1-11
1-8	Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation.....	1-12
4-1	x86-64 Oracle Linux 7 Minimum Operating System Requirements.....	4-12
4-2	x86-64 Oracle Linux 6 Minimum Operating System Requirements.....	4-13
4-3	x86-64 Red Hat Enterprise Linux 7 Minimum Operating System Requirements.....	4-16
4-4	x86-64 Red Hat Enterprise Linux 6 Minimum Operating System Requirements.....	4-17
4-5	x86-64 SUSE Linux Enterprise Server 12 Minimum Operating System Requirements..	4-18
4-6	Red Hat Enterprise Linux 7 Minimum Operating System Requirements.....	4-20
4-7	Red Hat Enterprise Linux 6 Minimum Operating System Requirements.....	4-21
4-8	SUSE Linux Enterprise Server 12 Minimum Operating System Requirements.....	4-22
4-9	Requirements for Programming Environments for Linux X86-64.....	4-26
5-1	Grid Naming Service Cluster Configuration Example.....	5-23
5-2	Manual Network Configuration Example.....	5-24
6-1	Installation Owner Resource Limit Recommended Ranges.....	6-25
7-1	Supported Storage Options for Oracle Grid Infrastructure.....	7-2
7-2	Platforms That Support Oracle ACFS and Oracle ADVM.....	7-4
8-1	Oracle ASM Disk Space Minimum Requirements for Oracle Database.....	8-7
8-2	Oracle ASM Disk Space Minimum Requirements for Oracle Database (non-CDB).....	8-7
8-3	Minimum Space Requirements for Oracle Domain Services Cluster with Four or Fewer Oracle Member Clusters.....	8-8
8-4	Minimum Space Requirements for Oracle Member Cluster.....	8-9
8-5	Minimum Space Requirements for Oracle Standalone Cluster.....	8-9
9-1	Oracle ASM Disk Group Redundancy Levels for Oracle Extended Clusters.....	9-6
11-1	Upgrade Checklist for Oracle Grid Infrastructure Installation.....	11-7
A-1	Response Files for Oracle Database and Oracle Grid Infrastructure.....	A-4
B-1	Device Name Formats Based on Disk Type.....	B-8
B-2	Disk Management Tasks Using ORACLEASM.....	B-10
B-3	Minimum Operating System Resource Parameter Settings.....	B-18
B-4	Commands to Display Kernel Parameter Values.....	B-19
C-1	Examples of OFA-Compliant Oracle Base Directory Names.....	C-4
C-2	Optimal Flexible Architecture Hierarchical File Path Examples.....	C-5

Preface

This guide explains how to configure a server in preparation for installing and configuring an Oracle Grid Infrastructure installation (Oracle Clusterware and Oracle Automatic Storage Management).

It also explains how to configure a server and storage in preparation for an Oracle Real Application Clusters (Oracle RAC) installation.

[Audience](#) (page xvii)

[Documentation Accessibility](#) (page xvii)

[Related Documentation](#) (page xviii)

[Conventions](#) (page xviii)

Audience

Oracle Grid Infrastructure Installation Guide provides configuration information for network and system administrators, and database installation information for database administrators (DBAs) who install and configure Oracle Clusterware and Oracle Automatic Storage Management in an Oracle Grid Infrastructure for a cluster installation.

For customers with specialized system roles who intend to install Oracle RAC, this book is intended to be used by system administrators, network administrators, or storage administrators to configure a system in preparation for an Oracle Grid Infrastructure for a cluster installation, and complete all configuration tasks that require operating system root privileges. When Oracle Grid Infrastructure installation and configuration is completed successfully, a system administrator should only need to provide configuration information and to grant access to the database administrator to run scripts as root during an Oracle RAC installation.

This guide assumes that you are familiar with Oracle Database concepts.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/>

[topic/lookup?ctx=acc&id=info](#) or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documentation

For more information, see the following Oracle resources:

Related Topics:

Oracle Real Application Clusters Installation Guide for Linux and UNIX

Oracle Database Installation Guide

Oracle Clusterware Administration and Deployment Guide

Oracle Real Application Clusters Administration and Deployment Guide

Oracle Database Concepts

Oracle Database New Features Guide

Oracle Database Licensing Information

Oracle Database Readme

Oracle Universal Installer User's Guide

Oracle Database Examples Installation Guide

Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems

Oracle Automatic Storage Management Administrator's Guide

Oracle Database Upgrade Guide

Oracle Database 2 Day DBA

Oracle Application Express Installation Guide

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Changes in This Release for Oracle Grid Infrastructure

This new release of Oracle Grid Infrastructure provides improvements to the installation process, performance, and automation.

[Changes in Oracle Grid Infrastructure 12c Release 2 \(12.2\)](#) (page xix)

[Changes in Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#) (page xxvii)

Changes in Oracle Grid Infrastructure 12c Release 2 (12.2)

The following are changes in *Oracle Grid Infrastructure Installation Guide* for Oracle Grid Infrastructure 12c Release 2 (12.2):

[New Features](#) (page xix)

[Deprecated Features](#) (page xxvi)

[Desupported Features](#) (page xxvi)

New Features

Following are the new features for Oracle Clusterware 12c Release 2 (12.2) and Oracle Automatic Storage Management 12c Release 2 (12.2):

New Features for Oracle Grid Infrastructure 12c Release 2 (12.2)

- Simplified Image-based Oracle Grid Infrastructure Installation

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure software is available as an image file for download and installation.

This feature greatly simplifies and enables quicker installation of Oracle Grid Infrastructure.

Note: You must extract the image software into the directory where you want your Grid home to be located, and then run the `gridSetup.sh` script to start Oracle Grid Infrastructure installation.

See Also: [About Image-Based Oracle Grid Infrastructure Installation](#) (page 9-2)

- Support for Oracle Domain Services Clusters and Oracle Member Clusters
Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure installer supports the option of deploying Oracle Domain Services Clusters and Oracle Member Clusters.

See Also: [Understanding Cluster Configuration Options](#) (page 9-3)

Note: Oracle Domain Services Clusters and Oracle Member Clusters are not supported on IBM: Linux on System z.

- Support for Oracle Extended Clusters
Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure installer supports the option of configuring cluster nodes in different locations as an Oracle Extended Cluster. An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

See Also: [Understanding Cluster Configuration Options](#) (page 9-3)

- Global Grid Infrastructure Management Repository
Oracle Grid Infrastructure deployment now supports a global off-cluster Grid Infrastructure Management Repository (GIMR). This repository is a multitenant database with a pluggable database (PDB) for the GIMR of each cluster. The global GIMR runs in an Oracle Domain Services Cluster. A global GIMR frees the local cluster from dedicating storage in its disk groups for this data and permitting longer term historical data storage for diagnostic and performance analysis.

See Also: [About the Grid Infrastructure Management Repository](#) (page 8-10)

Note: This feature is not supported on IBM: Linux on System z.

- Rapid Home Provisioning of Oracle Software
Rapid Home Provisioning enables you to create clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes. You can also provision Oracle Database on 11.2 clusters.

Rapid Home Provisioning leverages a new file system capability for separation of gold image software from the site-specific configuration changes. This separation ensures that the home path remains unchanged throughout updates. This feature combines the benefits of in-place and out-of-place patching. This capability is available with Oracle Grid Infrastructure 12c Release 2 (12.2).

See Also: [About Deploying Oracle Grid Infrastructure Using Rapid Home Provisioning](#) (page 9-31)

Note: Rapid Home Provisioning is not supported on IBM: Linux on System z.

- Parallel NFS Support in Oracle Direct NFS Client

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Direct NFS Client supports parallel NFS. Parallel NFS is an NFSv4.1 option that allows direct client access to file servers, enabling scalable distributed storage.

See Also: [About Direct NFS Client Mounts to NFS Storage Devices](#) (page 7-11)

- Direct NFS Dispatcher Support

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Direct NFS Client supports adding a dispatcher or I/O slave infrastructure. For very large database deployments running Oracle Direct NFS Client, this feature facilitates scaling of sockets and TCP connections to multi-path and clustered NFS storage.

See Also: [About Direct NFS Client Mounts to NFS Storage Devices](#) (page 7-11)

- Kerberos Authentication for Direct NFS

Oracle Database now supports Kerberos implementation with Direct NFS communication. This feature solves the problem of authentication, message integrity, and optional encryption over unsecured networks for data exchange between Oracle Database and NFS servers using Direct NFS protocols.

See Also: [Creating an oranfstab File for Direct NFS Client](#) (page 8-18) for information about setting up Kerberos authentication for Direct NFS Client

- Support for IPv6 Based IP Addresses for the Oracle Cluster Interconnect

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), you can use either IPv4 or IPv6 based IP addresses to configure cluster nodes on the private network. You can use more than one private network for the cluster.

- Shared Grid Naming Service (GNS) High Availability

Shared GNS High Availability provides high availability of lookup and other services to the clients by running multiple instances of GNS with primary and secondary roles.

See Also: *Oracle Clusterware Administration and Deployment Guide*

- Cluster Health Advisor

The Cluster Health Advisor provides system administrators and database administrators early warning of pending performance issues and root causes and corrective actions for Oracle RAC databases and cluster nodes. This advanced proactive diagnostic capability enhances availability and performance management.

Note: Cluster Health Advisor is not supported on IBM: Linux on System z.

- Enhancements to Cluster Verification Utility

Cluster Verification Utility (CVU) assists in the installation and configuration of Oracle Clusterware and Oracle Real Application Clusters (Oracle RAC). CVU performs a range of tests, covering all intermediate stages during the installation and configuration of a complete Oracle RAC stack. In this release, CVU provides several enhancements, such as information about the progress of each check and allowing you to specify an output format such as XML or HTML on request.

- Postinstallation Configuration of Oracle Software using the `-executeConfigTools` option

Starting with Oracle Database 12c Release 2 (12.2), you can perform postinstallation configuration of Oracle products by running the Oracle Database or Oracle Grid Infrastructure installer with the `-executeConfigTools` option. You can use the same response file created during installation to complete postinstallation configuration.

See Also: [Postinstallation Configuration Using Response File Created During Installation](#) (page A-10)

- Separation of Duty for Administering Oracle Real Application Clusters

Starting with Oracle Database 12c Release 2 (12.2), Oracle Database provides support for separation of duty best practices when administering Oracle Real Application Clusters (Oracle RAC) by introducing the SYSRAC administrative privilege for the clusterware agent. This feature removes the need to use the powerful SYSDBA administrative privilege for Oracle RAC.

SYSRAC, like SYSDG, SYSBACKUP and SYSKM, helps enforce separation of duties and reduce reliance on the use of SYSDBA on production systems. This administrative privilege is the default mode for connecting to the database by the clusterware agent on behalf of the Oracle RAC utilities such as `srvctl`.

- SCAN Listener Supports HTTP Protocol

Starting with Oracle Database 12c Release 2 (12.2), SCAN listener enables connections for the recovery server coming over HTTP to be redirected to different machines based on the load on the recovery server machines.

See Also: *Oracle Real Application Clusters Installation Guide*

- Oracle Real Application Clusters Reader Nodes

Oracle RAC Reader Nodes facilitate Oracle Flex Cluster architecture by allocating a set of read/write instances running Online Transaction Processing (OLTP) workloads and a set of read-only database instances across Hub Nodes and Leaf Nodes in the cluster. In this architecture, updates to the read-write instances are immediately propagated to the read-only instances on the Leaf Nodes, where they can be used for online reporting or instantaneous queries.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide*

- Service-Oriented Buffer Cache Access

Cluster-managed services are used to allocate workloads across various Oracle RAC database instances running in a cluster. These services are used to access database objects cached in the buffer caches of the respective database instances.

Service-oriented Buffer Cache Access optimization allows Oracle RAC to cache or pre-warm instances with data blocks for objects accessed through a service. This feature improves access time of Oracle RAC Database instances.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide*

- Server Weight-Based Node Eviction

Server weight-based node eviction acts as a tie-breaker mechanism in situations where Oracle Clusterware needs to evict a particular node or a group of nodes from a cluster, in which all nodes represent an equal choice for eviction. The server weight-based node eviction mechanism helps to identify the node or the group of nodes to be evicted based on additional information about the load on those servers. Two principle mechanisms, a system inherent automatic mechanism and a user input-based mechanism exist to provide respective guidance.

See Also: *Oracle Clusterware Administration and Deployment Guide*

- Load-Aware Resource Placement

Load-aware resource placement prevents overloading a server with more applications than the server is capable of running. The metrics used to determine whether an application can be started on a given server, either as part of the startup or as a result of a failover, are based on the anticipated resource consumption of the application as well as the capacity of the server in terms of CPU and memory.

New Features for Oracle Automatic Storage Management 12c Release 2 (12.2)

- Automatic Configuration of Oracle ASM Filter Driver

Oracle ASMFD simplifies the configuration and management of disk devices by eliminating the need to rebind disk devices used with Oracle ASM each time the system is restarted. The configuration for Oracle ASM Filter Driver (Oracle ASMFD) can now be enabled with a check box to be an automated process during Oracle Grid Infrastructure installation.

Note: Oracle ASMFD is supported on Linux x86-64 and Oracle Solaris operating systems.

See Also:

[About Oracle ASM with Oracle ASM Filter Driver](#) (page 8-14)

- Oracle IOServer

Oracle IOServer (IOS) provides connectivity for Oracle Database instances on nodes of member clusters that do not have connectivity to Oracle ASM managed disks. Oracle IOServer provides network-based file access to Oracle ASM disk groups.

See Also:

[About Oracle Flex ASM Clusters Networks](#) (page 5-19)

- Oracle ASM Flex Disk Groups and File Groups

Oracle ASM provides database-oriented storage management with flex disk groups. Flex disk groups support Oracle ASM file groups and quota groups.

See Also: *Oracle Automatic Storage Management Administrator's Guide*

- Oracle ACFS Snapshot-Based Replication

The new Oracle ACFS snapshot-based replication feature uses ACFS snapshot technology to transfer the differences between successive snapshots to the standby file system using standard ssh transport protocol. ACFS Snapshot-based replication is more efficient with higher performance, lower overhead, and ease of management.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Compression

Oracle ACFS provides file system compression functionality reducing storage requirement, resulting in lowering cost. Oracle ACFS compression is managed using the new `acfsutil compress` commands and updates to the `acfsutil info` command.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Defragger

Databases that share storage with snapshots or with the base of the file system can become fragmented under active online transaction processing (OLTP) workloads. This fragmentation can cause the location of the data in the volume to be discontinuous for sequential scans. Oracle ACFS automatically defragments these files in the background.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Support for 4K Sectors

Oracle ACFS supports I/O requests in multiples of 4K logical sector sizes as well as continued support for 512-byte logical sector size I/O requests. The `i 4096` option is provided with the `acfsformat` command on Windows and the `mkfs` command in Linux and Oracle Solaris environments.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Automatic Resize

Oracle ACFS provides an automatic resize option with the `acfsutil size` command. This command enables you to specify an increment by which an Oracle ACFS file system grows automatically if the amount of available free space in the file system falls below a specified amount. There is also an option to specify the maximum Oracle ACFS plugins support file content data collection. Both polling and interval based capture are supported with the file content collection. The output of the `acfsutil info fs` command displays the automatic resize increment and maximum amounts.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Metadata Acceleration

Oracle ACFS supports accelerator metadata storage. This support enables many critical Oracle ACFS metadata structures, including extent metadata, storage bitmaps, volume logs, and some snapshot metadata to be placed on accelerator storage.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Plugins for File Content Data Collection

Oracle ACFS plugins support file content data collection. Both polling and interval based capture are supported with the file content collection.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Sparse Files

Oracle ACFS provides support for sparse files. Oracle ACFS sparse files greatly benefit NFS client write operations which are commonly received out of order by the NFS server and the associated Oracle ACFS file system.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- Oracle ACFS Scrubbing Functionality

Oracle ACFS provides scrubbing functionality with the `acfsutil scrub` command to check for and report any inconsistencies in the metadata or file data.

See Also:

Oracle Automatic Storage Management Administrator's Guide

- High Availability Common Internet File System (HACIFS)
Oracle ACFS Common Internet File System (CIFS) features are enhanced to provide high availability for the exported file systems, with the Oracle ACFS NAS Maximum Availability eXtensions (NAS MAX) technology. High Availability Common Internet File System (HACIFS) and High Availability Network File System (HANFS) now both provide comprehensive Network Attach Storage solutions for Oracle ACFS customers.

See Also:

Oracle Automatic Storage Management Administrator's Guide

Deprecated Features

The following feature is deprecated in this release, and may be desupported in another release. See *Oracle Database Upgrade Guide* for a complete list of deprecated features in this release.

- Deprecation of `configToolAllCommands` script
The `configToolAllCommands` script runs in the response file mode to configure Oracle products after installation and uses a separate password response file. Starting with Oracle Database 12c Release 2 (12.2), the `configToolAllCommands` script is deprecated and is subject to desupport in a future release.
To perform postinstallation configuration of Oracle products, you can now run the Oracle Database or Oracle Grid Infrastructure installer with the `-executeConfigTools` option. You can use the same response file created during installation to complete postinstallation configuration.

Desupported Features

The following feature is desupported in this release. See *Oracle Database Upgrade Guide* for a complete list of features desupported in this release.

- Desupport of Direct File System Placement for Oracle Cluster Registry (OCR) and Voting Files
Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), the placement of Oracle Clusterware files: the Oracle Cluster Registry (OCR), and the Voting Files, directly on a shared file system is desupported in favor of having Oracle Clusterware files managed by Oracle Automatic Storage Management (Oracle ASM). You cannot place Oracle Clusterware files directly on a shared file system. If you need to use a supported shared file system, either a Network File System, or a shared cluster file system instead of native disk devices, then you must create Oracle ASM disks on supported network file systems that you plan to use for hosting Oracle Clusterware files before installing Oracle Grid Infrastructure. You can then use the Oracle ASM disks in an Oracle ASM disk group to manage Oracle Clusterware files.

If your Oracle Database files are stored on a shared file system, then you can continue to use shared file system storage for database files, instead of moving them to Oracle ASM storage.

Changes in Oracle Grid Infrastructure 12c Release 1 (12.1)

The following are changes in *Oracle Grid Infrastructure Installation Guide* for Oracle Grid Infrastructure 12c Release 1 (12.1):

[New Features for Oracle Grid Infrastructure 12c Release 1 \(12.1.0.2\)](#) (page xxvii)

[New Features for Oracle Grid Infrastructure 12c Release 1 \(12.1.0.1\)](#) (page xxviii)

[Deprecated Features](#) (page xxxi)

[Desupported Features](#) (page xxxi)

New Features for Oracle Grid Infrastructure 12c Release 1 (12.1.0.2)

- Automated Root Privilege Delegation to Scripts During Installation

You can continue to run scripts as root manually during installation, or you can enable OUI to run root scripts as needed during installation, using one of three methods: 1) providing the root password to the installer; 2) configuring Sudo access to the Oracle installation owner; 3) configuring PowerBroker access to the Oracle Installation owner.

- Oracle ASM Filter Driver

The Oracle ASM filter driver (Oracle ASMFD) is a kernel module that resides in the I/O path of the Oracle ASM disks. It is used by Oracle ASM to specify the parameters to validate all I/O requests.

The filter driver rejects any I/O requests that are not determined to be valid. This action eliminates accidental overwrites of Oracle ASM disks that would cause corruption in an Oracle ASM disk and files within the disk group. For example, the driver filters out all non-Oracle I/O requests that could cause accidental overwrites.

See *Oracle Automatic Storage Management Administrator's Guide* for more information about configuration and administration of Oracle ASMFD.

- Rapid Home Provisioning

Rapid Home Provisioning is a method of deploying software homes to nodes in a cloud computing environment from a single cluster where you store home images (called gold images) of Oracle software, such as databases, middleware, and applications. Rapid Home Provisioning Server (RHPS) clusters provide gold images to Rapid Home Provisioning Clients (RHPC).

See Also:

Oracle Clusterware Administration and Deployment Guide

- Cluster and Oracle RAC Diagnosability Tools Enhancements

The Trace File Analyzer (TFA) Collector is installed automatically with Oracle Grid Infrastructure installation. The Trace File Analyzer Collector is a diagnostic

collection utility to simplify diagnostic data collection on Oracle Grid Infrastructure and Oracle RAC systems.

See Also:

Oracle Clusterware Administration and Deployment Guide for information about using Trace File Analyzer Collector

- **Automatic Installation of Grid Infrastructure Management Repository**
The Grid Infrastructure Management Repository is automatically installed with Oracle Grid Infrastructure 12c Release 1 (12.1.0.2).
- **Oracle RAC Cache Fusion Accelerator**
Oracle RAC uses its Cache Fusion protocol and Global Cache Service (GCS) to provide fast, reliable, and efficient inter-instance data communication in an Oracle RAC cluster, so that the individual memory buffer caches of multiple instances can function as one global cache for the database. Using Cache Fusion provides a nearly linear scalability for most applications. This release includes accelerations to the Cache Fusion protocol that provide enhanced scalability for all applications.

New Features for Oracle Grid Infrastructure 12c Release 1 (12.1.0.1)

- **Automated Root Privilege Delegation to Scripts During Installation**
You can continue to run scripts as root manually during installation, or you can enable OUI to run root scripts as needed during installation, using one of three methods: 1) providing the root password to the installer; 2) configuring Sudo access to the Oracle installation owner; 3) configuring PowerBroker access to the Oracle Installation owner.
- **Database Upgrade Automation Using DBUA**
There are three areas that are being enhanced for upgrade ease-of-use. First, in the pre-upgrade phase, the existing manual steps are eliminated and give more explicit advice or even generate a fix-up script to find issues identified in the pre-upgrade phase. Second, in the post-upgrade phase, there is a post-upgrade health check that indicates that the upgrade was successful. Finally, partner documents (such as SAP) and major customer upgrade documents are used to further identify manual steps that may be automated and generalized to a wider customer base.

Automating the upgrade process provides major improvements in usability and ease-of-use. There is also better integration of database upgrade with Oracle Grid Infrastructure for a cluster and Oracle Enterprise Manager Cloud Control.
See Oracle Database Upgrade Guide.
- **DBCA Support for Multitenant Container Database and Pluggable Database Configurations**
Starting with Oracle Database 12c Release 1 (12.1), Oracle Database Configuration Assistant (DBCA) allows you to create either a multitenant container database (CDB) or a non-CDB. You can create the CDB with zero, one, or more pluggable databases (PDBs).

You can also create a CDB with one PDB during the database installation.

See *Oracle Database Administrator's Guide*.

- Enhancements to Cluster Health Monitor (CHM)

CHM has been enhanced to be more efficient to support Oracle Flex Clusters implementations. These enhancements ensure that Oracle Flex Clusters run smoothly while minimizing the required resources to monitor the stack.

See *Oracle Clusterware Administration and Deployment Guide*.

- Oracle Flex ASM Servers

Oracle Flex ASM enables the Oracle ASM instance to run on a separate physical server from the database servers. Many Oracle ASM instances can be clustered to support a large number of database clients.

Note that Oracle Flex ASM can apply to a collection of databases, each one a single instance but running in an Oracle Flex ASM Cluster.

See *Oracle Automatic Storage Management Administrator's Guide*.

- Oracle Flex Clusters

Oracle Flex Cluster is a new concept, which joins together a traditional closely coupled cluster with a modest node count with a large number of loosely coupled nodes. To support various configurations that can be established using this new concept, SRVCTL provides new commands and command options to ease the installation and configuration.

See *Oracle Clusterware Administration and Deployment Guide*.

- IPv6 Support for Public Networks

Oracle Clusterware 12c Release 1 (12.1) supports IPv6-based public IP and VIP addresses.

IPv6-based IP addresses have become the latest standard for the information technology infrastructure in today's data centers. With this release, Oracle RAC and Oracle Grid Infrastructure support this standard. You can configure cluster nodes during installation with either IPv4 or IPv6 addresses on the same network. Database clients can connect to either IPv4 or IPv6 addresses. The Single Client Access Name (SCAN) listener automatically redirects client connects to the appropriate database listener for the IP protocol of the client request.

See *Oracle Grid Infrastructure Installation Guide*.

- Multiprocess Multithreaded Oracle Database

Starting with Oracle Database 12c, Oracle Database may use operating system threads to allow resource sharing and reduce resource consumption.

See *Oracle Database Concepts*.

- Oracle ACFS Auditing and Support for Importing Auditing Data into Oracle Audit Vault Server

This feature provides auditing for Oracle ACFS security and encryption. In addition, this feature also generates an XML file containing Oracle ACFS audit trail data which can be imported by Oracle Audit Vault Server.

See *Oracle Automatic Storage Management Administrator's Guide*.

- Oracle Enterprise Manager Database Express 12c

Oracle Database 12c introduces Oracle Enterprise Manager Database Express, a web management product built into Oracle Database without any need for special installation or management. Using Oracle Enterprise Manager Database Express, you can perform administrative tasks such as managing user security, and managing database memory and storage. You can also view performance and status information about your database.

Note that starting with Oracle Database 12c, Oracle Enterprise Manager Database Control is deprecated.

See *Oracle Database 2 Day DBA*.

- Policy-Based Cluster Management and Administration

Oracle Grid Infrastructure allows running multiple applications in one cluster. Using a policy-based approach, the workload introduced by these applications can be allocated across the cluster using a policy. In addition, a policy set enables different policies to be applied to the cluster over time as required. Policy sets can be defined using a web-based interface or a command-line interface.

Hosting various workloads in the same cluster helps to consolidate the workloads into a shared infrastructure that provides high availability and scalability. Using a centralized policy-based approach allows for dynamic resource reallocation and prioritization as the demand changes.

See *Oracle Clusterware Administration and Deployment Guide*.

- Simplified Oracle Database Vault Installation

Starting with Oracle Database 12c, Oracle Database Vault is installed by default as part of the Oracle Database installation. However, you can configure, enable, or disable Oracle Database Vault after the Oracle Database installation, either using DBCA, or by running SQL statements.

See *Oracle Database Vault Administrator's Guide*.

- Support for Separation of Database Administration Duties

Oracle Database 12c provides support for separation of administrative duties for Oracle Database by introducing task-specific and least-privileged administrative privileges that do not require the SYSDBA administrative privilege. These new privileges are: SYSBACKUP for backup and recovery, SYSDG for Oracle Data Guard, and SYSKM for encryption key management.

See *Oracle Database Security Guide*.

- Unified Database Audit Configuration

Starting with Oracle Database 12c, you can create named audit policies. An audit policy contains a set of audit options, which is stored in the database as an object. The advantage of creating a named audit policy is that it reduces the number of commands that are required to create a database audit policy, and it simplifies the implementation of an audit configuration for security and compliance with conditional auditing. This new audit policy framework is included with the database installation.

See *Oracle Database Security Guide*.

Deprecated Features

The following features are deprecated in this release, and may be desupported in another release. See *Oracle Database Upgrade Guide* for a complete list of deprecated features in this release.

- Deprecation of single-letter SRVCTL command-line interface (CLI) options
All SRVCTL commands have been enhanced to accept full-word options instead of the single-letter options. All new SRVCTL command options added in this release support full-word options, only, and do not have single-letter equivalents. The use of single-letter options with SRVCTL commands might be desupported in a future release.
- Change for Standalone Deinstallation Tool
The deinstallation tool is now integrated with the database installation media.
- Deprecation of `-cleanupOBase`
The `-cleanupOBase` flag of the deinstallation tool is deprecated in this release. There is no replacement for this flag.
- Oracle Enterprise Manager Database Control is replaced by Oracle Enterprise Manager Database Express.
- The `deinstall` standalone utility is replaced with a `deinstall` option using Oracle Universal Installer (OUI).

Desupported Features

The following features are no longer supported by Oracle. See *Oracle Database Upgrade Guide* for a complete list of features desupported in this release.

- Oracle Enterprise Manager Database Control
- `CLEANUP_ORACLE_BASE` property removed and does not support an Oracle base removal during silent or response file mode deinstalls.

Oracle Grid Infrastructure Installation Checklist

Use checklists to plan and carry out Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management).

Oracle recommends that you use checklists as part of your installation planning process. Using this checklist can help you to confirm that your server hardware and configuration meet minimum requirements for this release, and can help you to ensure you carry out a successful installation.

[Server Hardware Checklist for Oracle Grid Infrastructure](#) (page 1-1)

Review server hardware requirements for Oracle Grid Infrastructure installation.

[Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC](#) (page 1-2)

Review the checklist for operating system requirements for Oracle Grid Infrastructure installation.

[Server Configuration Checklist for Oracle Grid Infrastructure](#) (page 1-4)

Use this checklist to check minimum server configuration requirements for Oracle Grid Infrastructure installations.

[Network Checklist for Oracle Grid Infrastructure](#) (page 1-5)

Review this network checklist for Oracle Grid Infrastructure installation to ensure that you have required hardware, names, and addresses for the cluster.

[User Environment Configuration Checklist for Oracle Grid Infrastructure](#) (page 1-7)

Use this checklist to plan operating system users, groups, and environments for Oracle Grid Infrastructure installation.

[Storage Checklist for Oracle Grid Infrastructure](#) (page 1-9)

Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.

[Cluster Deployment Checklist for Oracle Grid Infrastructure](#) (page 1-11)

Review the checklist for planning your cluster deployment Oracle Grid Infrastructure installation.

[Installer Planning Checklist for Oracle Grid Infrastructure](#) (page 1-12)

Review the checklist for planning your Oracle Grid Infrastructure installation before starting Oracle Universal Installer.

1.1 Server Hardware Checklist for Oracle Grid Infrastructure

Review server hardware requirements for Oracle Grid Infrastructure installation.

Table 1-1 Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Server make and architecture	Confirm that server makes, models, core architecture, and host bus adaptors (HBA) are supported to run with Oracle Grid Infrastructure and Oracle RAC.
Runlevel	3 or 5
Server Display Cards	At least 1024 x 768 display resolution for Oracle Universal Installer. Confirm display monitor.
Minimum Random Access Memory (RAM)	At least 8 GB RAM for Oracle Grid Infrastructure installations.
Intelligent Platform Management Interface (IPMI)	IPMI cards installed and configured, with IPMI administrator account information available to the person running the installation. Ensure baseboard management controller (BMC) interfaces are configured, and have an administration account username and password to provide when prompted during installation.

1.2 Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC

Review the checklist for operating system requirements for Oracle Grid Infrastructure installation.

Table 1-2 Operating System General Checklist for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Operating system general requirements	<p>OpenSSH installed manually, if you do not have it installed already as part of a default Linux installation.</p> <p>Review the system requirements section for a list of minimum package requirements. Use the same operating system kernel and packages on each cluster member node.</p> <ul style="list-style-type: none"> Oracle Linux 7 with the Unbreakable Enterprise Kernel 3: 3.8.13-35.3.1.el7uek.x86_64 or later Oracle Linux 7.2 with the Unbreakable Enterprise Kernel 4: 4.1.12-32.2.3.el7uek.x86_64 or later Oracle Linux 7 with the Red Hat Compatible kernel: 3.10.0-123.el7.x86_64 or later Oracle Linux 6.4 with the Unbreakable Enterprise Kernel 2: 2.6.39-400.211.1.el6uek.x86_64 or later Oracle Linux 6.6 with the Unbreakable Enterprise Kernel 3: 3.8.13-44.1.1.el6uek.x86_64 or later Oracle Linux 6.8 with the Unbreakable Enterprise Kernel 4: 4.1.12-37.6.2.el6uek.x86_64 or later Oracle Linux 6.4 with the Red Hat Compatible kernel: 2.6.32-358.el6.x86_64 or later Red Hat Enterprise Linux 7: 3.10.0-123.el7.x86_64 or later Red Hat Enterprise Linux 6.4: 2.6.32-358.el6.x86_64 or later SUSE Linux Enterprise Server 12 SP1: 3.12.49-11.1 or later <p>Review the system requirements section for a list of minimum package requirements.</p>
IBM: Linux on System z operating system requirements	<p>The following IBM: Linux on System z kernels are supported:</p> <p>Red Hat Enterprise Linux 7.2: 3.10.0-327.el7.s390x or later</p> <p>Red Hat Enterprise Linux 6.6: 2.6.32-504.el6.s390x or later</p> <p>SUSE Linux Enterprise Server 12 SP1: 3.12.49-11-default s390x or later</p> <p>Review the system requirements section for a list of minimum package requirements.</p>
Oracle Preinstallation RPM for Oracle Linux	<p>If you use Oracle Linux, then Oracle recommends that you run an Oracle preinstallation RPM for your Linux release to configure your operating system for Oracle Database and Oracle Grid Infrastructure installations.</p>
Oracle RPM Checker utility for IBM: Linux on System z	<p>Oracle recommends that you use the Oracle RPM Checker utility to verify that you have the required Red Hat Enterprise Linux or SUSE packages installed on your IBM: Linux on System z operating system before you start the Oracle Database or Oracle Grid Infrastructure installation.</p>
Disable Transparent HugePages	<p>Oracle recommends that you disable Transparent HugePages and use standard HugePages for enhanced performance.</p>

1.3 Server Configuration Checklist for Oracle Grid Infrastructure

Use this checklist to check minimum server configuration requirements for Oracle Grid Infrastructure installations.

Table 1-3 Server Configuration Checklist for Oracle Grid Infrastructure

Check	Task
Disk space allocated to the temporary file system	At least 1 GB of space in the temporary directory. Oracle recommends 2 GB or more.
Swap space allocation relative to RAM	Between 4 GB and 16 GB: Equal to RAM More than 16 GB: 16 GB Note: If you enable HugePages for your Linux servers, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.
HugePages memory allocation	Allocate memory to HugePages large enough for the System Global Areas (SGA) of all databases planned to run on the cluster, and to accommodate the System Global Area for the Grid Infrastructure Management Repository.
Mount point paths for the software binaries	Oracle recommends that you create an Optimal Flexible Architecture configuration.
Ensure that the Oracle home (the Oracle home path you select for Oracle Database) uses only ASCII characters	The ASCII character restriction includes installation owner user names, which are used as a default for some home paths, as well as other directory names you may select for paths.
Set locale (if needed)	Specify the language and the territory, or locale, in which you want to use Oracle components. A locale is a linguistic and cultural environment in which a system or program is running. NLS (National Language Support) parameters determine the locale-specific behavior on both servers and clients. The locale setting of a component determines the language of the user interface of the component, and the globalization behavior, such as date and number formatting.
Set Network Time Protocol for Cluster Time Synchronization	Oracle Clusterware requires the same time zone environment variable setting on all cluster nodes. Ensure that you set the time zone synchronization across all cluster nodes using either an operating system configured network time protocol (NTP) or Oracle Cluster Time Synchronization Service.
Check Shared Memory File System Mount	By default, your operating system includes an entry in <code>/etc/fstab</code> to mount <code>/dev/shm</code> . However, if your Cluster Verification Utility (CVU) or Oracle Universal Installer (OUI) checks fail, ensure that the <code>/dev/shm</code> mount area is of type <code>tmpfs</code> and is mounted with the following options: <ul style="list-style-type: none"> <code>rw</code> and <code>exec</code> permissions set on it Without <code>noexec</code> or <code>nosuid</code> set on it Note: These options might not be listed as they are usually set as the default permissions by your operating system.

Related Topics:[Optimal Flexible Architecture](#) (page C-1)

Oracle Optimal Flexible Architecture (OFA) rules are a set of configuration guidelines created to ensure well-organized Oracle installations, which simplifies administration, support and maintenance.

1.4 Network Checklist for Oracle Grid Infrastructure

Review this network checklist for Oracle Grid Infrastructure installation to ensure that you have required hardware, names, and addresses for the cluster.

During installation, you designate interfaces for use as public, private, or Oracle ASM interfaces. You can also designate interfaces that are in use for other purposes, such as a network file system, and not available for Oracle Grid Infrastructure use.

If you use a third-party cluster software, then the public host name information is obtained from that software.

Table 1-4 Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Public network hardware	<ul style="list-style-type: none"> Public network switch (redundant switches recommended) connected to a public gateway and to the public interface ports for each cluster member node. Ethernet interface card (redundant network cards recommended, bonded as one Ethernet port name). The switches and network interfaces must be at least 1 GbE. The network protocol is Transmission Control Protocol (TCP) and Internet Protocol (IP).
Private network hardware for the interconnect	<ul style="list-style-type: none"> Private dedicated network switches (redundant switches recommended), connected to the private interface ports for each cluster member node. <p>Note: If you have more than one private network interface card for each server, then Oracle Clusterware automatically associates these interfaces for the private network using Grid Interprocess Communication (GIPC) and Grid Infrastructure Redundant Interconnect, also known as Cluster High Availability IP (HAIP).</p> <ul style="list-style-type: none"> The switches and network interface adapters must be at least 1 GbE. The interconnect must support the user datagram protocol (UDP). Jumbo Frames (Ethernet frames greater than 1500 bits) are not an IEEE standard, but can reduce UDP overhead if properly configured. Oracle recommends the use of Jumbo Frames for interconnects. However, be aware that you must load-test your system, and ensure that they are enabled throughout the stack.
Oracle Flex ASM Network Hardware	Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. Each network can be classified PUBLIC or PRIVATE+ASM or PRIVATE or ASM. Oracle ASM networks use the TCP protocol.

Table 1-4 (Cont.) Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Cluster Names and Addresses	<p>Determine and configure the following names and addresses for the cluster:</p> <ul style="list-style-type: none"> • Cluster name: Decide a name for the cluster, and be prepared to enter it during installation. The cluster name should have the following characteristics: <ul style="list-style-type: none"> Globally unique across all hosts, even across different DNS domains. At least one character long and less than or equal to 15 characters long. Consist of the same character set used for host names, in accordance with RFC 1123: Hyphens (-), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9). If you use third-party vendor clusterware, then Oracle recommends that you use the vendor cluster name. • Grid Naming Service Virtual IP Address (GNS VIP): If you plan to use GNS, then configure a GNS name and fixed address in DNS for the GNS VIP, and configure a subdomain on your DNS delegated to the GNS VIP for resolution of cluster addresses. GNS domain delegation is mandatory with dynamic public networks (DHCP, autoconfiguration). • Single Client Access Name (SCAN) and addresses <p>Using Grid Naming Service Resolution: Do not configure SCAN names and addresses in your DNS. SCAN names are managed by GNS.</p> <p>Using Manual Configuration and DNS resolution: Configure a SCAN name to resolve to three addresses on the domain name service (DNS).</p>

Table 1-4 (Cont.) Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Hub Node Public, Private and Virtual IP names and Addresses	<p>If you are not using GNS, then configure the following for each Hub Node:</p> <ul style="list-style-type: none"> • Public node name and address, configured in the DNS and in <code>/etc/hosts</code> (for example, <code>node1.example.com</code>, address 192.0.2.10). The public node name should be the primary host name of each node, which is the name displayed by the <code>hostname</code> command. • Private node address, configured on the private interface for each node. <p>The private subnet that the private interfaces use must connect all the nodes you intend to have as cluster members. Oracle recommends that the network you select for the private network uses an address range defined as private by RFC 1918.</p> <ul style="list-style-type: none"> • Public node virtual IP name and address (for example, <code>node1-vip.example.com</code>, address 192.0.2.11). <p>If you are not using dynamic networks with GNS and subdomain delegation, then determine a virtual host name for each node. A virtual host name is a public node name that is used to reroute client requests sent to the node if the node is down. Oracle Database uses VIPs for client-to-database connections, so the VIP address must be publicly accessible. Oracle recommends that you provide a name in the format <code>hostname-vip</code>. For example: <code>myclstr2-vip</code>.</p> <p>If you are using GNS, then you can also configure Leaf Nodes on both public and private networks, during installation. Leaf Nodes on public networks do not use Oracle Clusterware services such as the public network resources and VIPs, or run listeners. After installation, you can configure network resources and listeners for the Leaf Nodes using <code>SRVCTL</code> commands.</p>

Related Topics:[Understanding Oracle Flex Clusters](#) (page 5-18)

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure cluster configurations are Oracle Flex Clusters deployments.

[About Oracle Flex ASM Clusters Networks](#) (page 5-19)

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), as part of an Oracle Flex Cluster installation, Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services.

[General Requirements for Oracle Flex Cluster Configuration](#) (page 5-21)

Review this information about network requirements for Oracle Flex Cluster configuration.

1.5 User Environment Configuration Checklist for Oracle Grid Infrastructure

Use this checklist to plan operating system users, groups, and environments for Oracle Grid Infrastructure installation.

Table 1-5 User Environment Configuration for Oracle Grid Infrastructure

Check	Task
Review Oracle Inventory (oraInventory) and OINSTALL Group Requirements	<p>The Oracle Inventory directory is the central inventory of Oracle software installed on your system. It should be the primary group for all Oracle software installation owners. Users who have the Oracle Inventory group as their primary group are granted the OINSTALL privilege to read and write to the central inventory.</p> <ul style="list-style-type: none"> • If you have an existing installation, then OUI detects the existing oraInventory directory from the <code>/etc/oraInst.loc</code> file, and uses this location. • If you are installing Oracle software for the first time, then OUI creates an Oracle base and central inventory, and creates an Oracle inventory using information in the following priority: <ul style="list-style-type: none"> – In the path indicated in the ORACLE_BASE environment variable set for the installation owner user account – In an Optimal Flexible Architecture (OFA) path (<code>u[01–99]/app/owner</code> where <i>owner</i> is the name of the user account running the installation), if that user account has permissions to write to that path – In the user home directory, in the path <code>/app/owner</code>, where <i>owner</i> is the name of the user account running the installation <p>Ensure that the group designated as the OINSTALL group is available as the primary group for all planned Oracle software installation owners.</p>
Create operating system groups and users for standard or role-allocated system privileges	<p>Create operating system groups and users depending on your security requirements, as described in this installation guide.</p> <p>Set resource limits settings and other requirements for Oracle software installation owners.</p> <p>Group and user names must use only ASCII characters.</p>
Unset Oracle Software Environment Variables	<p>If you have an existing Oracle software installation, and you are using the same user to install this installation, then unset the following environment variables: <code>\$ORACLE_HOME</code>; <code>\$ORA_NLS10</code>; <code>\$TNS_ADMIN</code>.</p> <p>If you have set <code>\$ORA_CRS_HOME</code> as an environment variable, then unset it before starting an installation or upgrade. Do not use <code>\$ORA_CRS_HOME</code> as a user environment variable, except as directed by Oracle Support.</p>
Configure the Oracle Software Owner Environment	<p>Configure the environment of the oracle or grid user by performing the following tasks:</p> <ul style="list-style-type: none"> • Set the default file mode creation mask (umask) to 022 in the shell startup file. • Set the DISPLAY environment variable.

Table 1-5 (Cont.) User Environment Configuration for Oracle Grid Infrastructure

Check	Task
Determine root privilege delegation option for installation	<p>During installation, you are asked to run configuration scripts as the root user. You can either run these scripts manually as root when prompted, or during installation you can provide configuration information and passwords using a root privilege delegation option.</p> <p>To run root scripts automatically, select Automatically run configuration scripts, during installation. To use the automatic configuration option, the root user credentials for all cluster member nodes must use the same password.</p> <ul style="list-style-type: none"> • Use root user credentials Provide the superuser password for cluster member node servers. • Use sudo sudo is a UNIX and Linux utility that allows members of the <code>sudoers</code> list privileges to run individual commands as root. Provide the user name and password of an operating system user that is a member of <code>sudoers</code>, and is authorized to run <code>sudo</code> on each cluster member node. To enable <code>sudo</code>, have a system administrator with the appropriate privileges configure a user that is a member of the <code>sudoers</code> list, and provide the user name and password when prompted during installation.

1.6 Storage Checklist for Oracle Grid Infrastructure

Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.

Table 1-6 Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
Minimum disk space (local or shared) for Oracle Grid Infrastructure Software	<ul style="list-style-type: none"> • At least 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home). Oracle recommends that you allocate 100 GB to allow additional space for patches. At least 9 GB for Oracle Database Enterprise Edition • Allocate additional storage space as per your cluster configuration, as described in <i>Oracle Clusterware Storage Space Requirements</i>.

Table 1-6 (Cont.) Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
Select Oracle ASM Storage Options	<p>During installation, based on the cluster configuration, you are asked to provide Oracle ASM storage paths for the Oracle Clusterware files. These path locations must be writable by the Oracle Grid Infrastructure installation owner (Grid user). These locations must be shared across all nodes of the cluster on Oracle ASM because the files in the Oracle ASM disk group created during installation must be available to all cluster member nodes.</p> <ul style="list-style-type: none"> For Oracle Standalone Cluster deployment, shared storage, either Oracle ASM or Oracle ASM on NFS, is locally mounted on each of the Hub Nodes. For Oracle Domain Services Cluster deployment, Oracle ASM storage is shared across all nodes, and is available to Oracle Member Clusters. Oracle Member Cluster for Oracle Databases can either use storage services from the Oracle Domain Services Cluster or local Oracle ASM storage shared across all the nodes. Oracle Member Cluster for Applications always use storage services from the Oracle Domain Services Cluster. Before installing Oracle Member Cluster, create a Member Cluster Manifest file that specifies the storage details. <p>Voting files are files that Oracle Clusterware uses to verify cluster node membership and status. Oracle Cluster Registry files (OCR) contain cluster and database configuration information for Oracle Clusterware.</p>
Select Grid Infrastructure Management Repository (GIMR) Storage Option	<p>Depending on the type of cluster you are installing, you can choose to either host the Grid Infrastructure Management Repository (GIMR) for a cluster on the same cluster or on a remote cluster.</p> <p>For Oracle Standalone Cluster deployment, you can specify the same or separate Oracle ASM disk group for the GIMR.</p> <p>For Oracle Domain Services Cluster deployment, the GIMR must be configured on a separate Oracle ASM disk group.</p> <p>Oracle Member Clusters use the remote GIMR of the Oracle Domain Services Cluster. You must specify the GIMR details when you create the Member Cluster Manifest file before installation.</p>

Related Topics:[Oracle Clusterware Storage Space Requirements](#) (page 8-7)

Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files, and installing the starter database, for various Oracle Cluster deployments.

Related Topics:[About the Grid Infrastructure Management Repository](#) (page 8-10)

Every Oracle Standalone Cluster and Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR), or the Management Database (MGMTDB).

[Creating Member Cluster Manifest File for Oracle Member Clusters](#) (page 8-21)
Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Rapid Home Provisioning configuration.

[About Oracle Standalone Clusters](#) (page 9-3)
An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

[About Oracle Cluster Domain and Oracle Domain Services Cluster](#) (page 9-3)
An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.

[About Oracle Member Clusters](#) (page 9-5)
Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications.

[Understanding Cluster Configuration Options](#) (page 9-3)
Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 12c Release 2.

1.7 Cluster Deployment Checklist for Oracle Grid Infrastructure

Review the checklist for planning your cluster deployment Oracle Grid Infrastructure installation.

Table 1-7 Oracle Grid Infrastructure Cluster Deployment Checklist

Check	Task
Configure an Oracle Cluster that hosts all Oracle Grid Infrastructure services and Oracle ASM locally and accesses storage directly	<p>Deploy an Oracle Standalone Cluster.</p> <p>Use the Oracle Extended Cluster option to extend an Oracle RAC cluster across two, or more, separate sites, each equipped with its own storage.</p>
Configure an Oracle Cluster Domain to standardize, centralize, and optimize your Oracle Real Application Clusters (Oracle RAC) deployment	<p>Deploy an Oracle Domain Services Cluster.</p> <p>To run Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database instances, deploy Oracle Member Cluster for Oracle Databases.</p> <p>To run highly-available software applications, deploy Oracle Member Cluster for Applications.</p>

Related Topics:

[Installing Oracle Standalone Cluster](#) (page 9-7)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

[Installing Oracle Domain Services Cluster](#) (page 9-14)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.

[Installing Oracle Member Clusters](#) (page 9-21)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

Related Topics:

[Understanding Cluster Configuration Options](#) (page 9-3)

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 12c Release 2.

1.8 Installer Planning Checklist for Oracle Grid Infrastructure

Review the checklist for planning your Oracle Grid Infrastructure installation before starting Oracle Universal Installer.

Table 1-8 Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation

Check	Task
Read the Release Notes	Review release notes for your platform, which are available for your release at the following URL: http://www.oracle.com/technetwork/indexes/documentation/index.html
Review the Licensing Information	You are permitted to use only those components in the Oracle Database media pack for which you have purchased licenses. Refer to <i>Oracle Database Licensing Information</i> for more information about licenses.
Run OUI with CVU and use fixup scripts	Oracle Universal Installer is fully integrated with Cluster Verification Utility (CVU), automating many CVU prerequisite checks. Oracle Universal Installer runs all prerequisite checks and creates fixup scripts when you run the installer. You can run OUI up to the Summary screen without starting the installation. You can also run CVU commands manually to check system readiness. For more information, see: <i>Oracle Clusterware Administration and Deployment Guide</i>
Download and run ORAchk for runtime and upgrade checks, or runtime health checks	The ORAchk utility provides system checks that can help to prevent issues after installation. These checks include kernel requirements, operating system resource allocations, and other system requirements. Use the ORAchk Upgrade Readiness Assessment to obtain an automated upgrade-specific system health check for upgrades. For example: <pre>./orachk -u -o pre</pre> The ORAchk Upgrade Readiness Assessment automates many of the manual pre- and post-upgrade checks described in Oracle upgrade documentation. ORAchk is supported on Windows platforms in a Cygwin environment only. For more information refer to the following URL: https://support.oracle.com/rs?type=doc&id=1268927.1

Table 1-8 (Cont.) Oracle Universal Installer Checklist for Oracle Grid Infrastructure Installation

Check	Task
Ensure cron jobs do not run during installation	If the installer is running when daily cron jobs start, then you may encounter unexplained installation problems if your cron job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily cron jobs are run, or disable daily cron jobs that perform cleanup until after the installation is completed.
Obtain Your My Oracle Support account information.	During installation, you require a My Oracle Support user name and password to configure security updates, download software updates, and other installation tasks. You can register for My Oracle Support at the following URL: https://support.oracle.com/
Check running Oracle processes, and shut down processes if necessary.	<ul style="list-style-type: none"> • On a node with a standalone database not using Oracle ASM: You do not need to shut down the database while you install Oracle Grid Infrastructure. • On a node with a standalone Oracle Database using Oracle ASM: Stop the existing Oracle ASM instances. The Oracle ASM instances are restarted during installation. • On an Oracle RAC Database node: This installation requires an upgrade of Oracle Clusterware, as Oracle Clusterware is required to run Oracle RAC. As part of the upgrade, you must shut down the database one node at a time as the rolling upgrade proceeds from node to node.

Checking and Configuring Server Hardware for Oracle Grid Infrastructure

Verify that servers where you install Oracle Grid Infrastructure meet the minimum requirements for installation

This section provides minimum server requirements to complete installation of Oracle Grid Infrastructure. It does not provide system resource guidelines, or other tuning guidelines for particular workloads.

[Logging In to a Remote System Using X Window System](#) (page 2-1)

Use this procedure to run Oracle Universal Installer (OUI) by logging on to a remote system where the runtime setting prohibits logging in directly to a graphical user interface (GUI).

[Checking Server Hardware and Memory Configuration](#) (page 2-2)

Use this procedure to gather information about your server configuration.

2.1 Logging In to a Remote System Using X Window System

Use this procedure to run Oracle Universal Installer (OUI) by logging on to a remote system where the runtime setting prohibits logging in directly to a graphical user interface (GUI).

OUI is a graphical user interface (GUI) application. On servers where the runtime settings prevent GUI applications from running, you can redirect the GUI display to a client system connecting to the server.

Note:

If you log in as another user (for example, `oracle` or `grid`), then repeat this procedure for that user as well.

1. Start an X Window System session. If you are using an X Window System terminal emulator from a PC or similar system, then you may need to configure security settings to permit remote hosts to display X applications on your local system.
2. Enter a command using the following syntax to enable remote hosts to display X applications on the local X server:

```
# xhost + RemoteHost
```

RemoteHost is the fully qualified remote host name. For example:

```
# xhost + somehost.example.com
somehost.example.com being added to the access control list
```

3. If you are not installing the software on the local system, then use the `ssh` command to connect to the system where you want to install the software:

```
# ssh -Y RemoteHost
```

RemoteHost is the fully qualified remote host name. The `-Y` flag ("yes") enables remote X11 clients to have full access to the original X11 display. For example:

```
# ssh -Y somehost.example.com
```

4. If you are not logged in as the `root` user, and you are performing configuration steps that require `root` user privileges, then switch the user to `root`.

Note:

For more information about remote login using X Window System, refer to your X server documentation, or contact your X server vendor or system administrator. Depending on the X server software that you are using, you may have to complete the tasks in a different order.

2.2 Checking Server Hardware and Memory Configuration

Use this procedure to gather information about your server configuration.

1. Use the following command to determine physical RAM size on the server:

```
# grep MemTotal /proc/meminfo
```

If the size of the physical RAM installed in the system is less than the required size, then you must install more memory before continuing.

2. Determine the size of the configured swap space:

```
# grep SwapTotal /proc/meminfo
```

If necessary, see your operating system documentation for information about how to configure additional swap space.

3. Determine the amount of space available in the `/tmp` directory:

```
# df -h /tmp
```

If the free space available in the `/tmp` directory is less than what is required, then complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to meet the disk space requirement.
- When you set the Oracle user's environment, also set the `TMP` and `TMPDIR` environment variables to the directory you want to use instead of `/tmp`.

4. Determine the amount of free RAM and disk swap space on the system:

```
# free
```

5. Determine if the system architecture can run the software:

```
# uname -m
```


Verify that the processor architecture matches the Oracle software release to install. For example, you should see the following for a x86-64 bit system:

```
x86_64
```

If you do not see the expected output, then you cannot install the software on this system.

6. Verify that shared memory (`/dev/shm`) is mounted properly with sufficient size:

```
df -h /dev/shm
```

The `df -h` command displays the filesystem on which `/dev/shm` is mounted, and also displays in GB the total size and free size of shared memory.

Automatically Configuring Oracle Linux with Oracle Preinstallation RPM

Use Oracle Linux RPMs to simplify operating system configuration in preparation for Oracle software installations.

Oracle recommends that you install Oracle Linux 7 or Oracle Linux 6 and use Oracle RPMs to configure your operating systems for Oracle Database and Oracle Grid Infrastructure installations. For Oracle Linux 6 and Oracle Linux 7, run the Oracle Preinstallation RPM.

[Overview of Oracle Linux Configuration with Oracle RPMs](#) (page 3-1)

Use Oracle Linux RPMs to simplify operating system configuration, and to ensure that you have required kernel packages.

[Installing the Oracle Preinstallation RPM with ULN Support](#) (page 3-2)

Use this procedure to subscribe to Unbreakable Linux Network (ULN) Oracle Linux channels for your Oracle software.

[Installing the Oracle Preinstallation RPM From Unbreakable Linux Network](#) (page 3-3)

Use this procedure to subscribe to Oracle Linux channels, and to add the Oracle Linux channel that distributes the Oracle Preinstallation RPM:

[Installing a New Oracle Linux Installation from DVDs or Images](#) (page 3-4)

Use this procedure to install a new Oracle Linux installation and to perform system configuration with the Oracle Preinstallation RPM:

[Installing Oracle Linux with Oracle Linux Yum Server Support](#) (page 3-5)

Use this procedure to install Oracle Linux and configure your Linux installation for security errata or bug fix updates using the Oracle Linux yum server.

[Configuring Oracle Ksplice to Perform Oracle Linux Updates](#) (page 3-6)

You can configure Oracle Ksplice to perform zero downtime updates for your operating system if you have an Oracle Linux Premier Support subscription and an access key, which is available on Unbreakable Linux Network (ULN).

[Configure Additional Operating System Features](#) (page 3-6)

Oracle recommends that you configure your operating system before starting installation with additional features, such as IPMI or additional programming environments.

3.1 Overview of Oracle Linux Configuration with Oracle RPMs

Use Oracle Linux RPMs to simplify operating system configuration, and to ensure that you have required kernel packages.

Oracle recommends that you install Oracle Linux 6 or Oracle Linux 7 and use Oracle RPMs to configure your operating systems for Oracle Database and Oracle Grid Infrastructure installations.

The Oracle RPMs for your Oracle Linux distributions and Oracle RDBMS releases automatically install any additional packages needed for installing Oracle Grid Infrastructure and Oracle Database, and configure your server operating system automatically, including setting kernel parameters and other basic operating system requirements for installation. For more information about what the Oracle RPMs do, refer to the following URL:

<https://linux.oracle.com>

Configuring a server using Oracle Linux and the Oracle Preinstallation RPM consists of the following steps:

1. Install Oracle Linux.
2. Register your Linux distribution with Oracle Unbreakable Linux Network (ULN) or download and configure the yum repository for your system using the Oracle Linux yum server for your Oracle Linux release.
3. Install the Oracle Preinstallation RPM with the RPM for your Oracle Grid Infrastructure and Oracle Database releases, and update your Linux release.
4. Create role-allocated groups and users with identical names and ID numbers.
5. Complete network interface configuration for each cluster node candidate.
6. Complete system configuration for shared storage access as required for each standard or Core node cluster candidate.

After these steps are complete, you can proceed to install Oracle Grid Infrastructure and Oracle RAC.

3.2 Installing the Oracle Preinstallation RPM with ULN Support

Use this procedure to subscribe to Unbreakable Linux Network (ULN) Oracle Linux channels for your Oracle software.

To obtain Unbreakable Linux Network (ULN) support, subscribe to Oracle Linux channels, and to add the Oracle Linux channel that distributes the Oracle Preinstallation RPM:

1. Register your server with Unbreakable Linux Network (ULN). By default, you are registered for the Oracle Linux Latest channel for your operating system and hardware.
2. Log in to Unbreakable Linux Network:
<https://linux.oracle.com>
3. Click the **Systems** tab, and in the System Profiles list, select a registered server. The System Details window opens and displays the subscriptions for the server.
4. Click **Manage Subscriptions** to open the System Summary window.
5. From the **Available Channels** list, select the latest and update patch channels corresponding to your Oracle Linux distribution. For example, if your distribution is Oracle Linux 7 for x86_64, then select **Oracle Linux 7 Latest (x86_64)**.

6. Click **Subscribe**.
7. Start a terminal session and enter the following command as `root`, depending on your platform. For example:

- Oracle Linux 6 and Oracle Linux 7:

```
# yum install oracle-database-server-12cR2-preinstall
```

Note: Use the `-y` option if you want yum to skip the package confirmation prompt.

You should see output indicating that you have subscribed to the Oracle Linux channel, and that packages are being installed.

Oracle Linux automatically creates a standard (not role-allocated) Oracle installation owner and groups, and sets up other kernel configuration settings as required for Oracle installations.

8. Check the RPM log file to review the system configuration changes. For example, on Oracle Linux 7:

```
/var/log/oracle-database-server-12cR2-preinstall/backup/timestamp/orakernel.log
```

9. Repeat steps 1 through 8 on all other servers in your cluster.

3.3 Installing the Oracle Preinstallation RPM From Unbreakable Linux Network

Use this procedure to subscribe to Oracle Linux channels, and to add the Oracle Linux channel that distributes the Oracle Preinstallation RPM:

1. Complete a default Oracle Linux workstation installation.

You can download Oracle Linux from the Oracle Software Delivery Cloud:

<https://edelivery.oracle.com/linux>

2. Register your server with Unbreakable Linux Network (ULN). By default, you are registered for the Oracle Linux Latest channel for your operating system and hardware.
3. Log in to Unbreakable Linux Network:

<https://linux.oracle.com>
4. Click the **Systems** tab, and in the System Profiles list, select a registered server. The System Details window opens and displays the subscriptions for the server.
5. Click **Manage Subscriptions** to open the System Summary window.
6. From the **Available Channels** list, select the Latest and update patch channels corresponding to your Oracle Linux distribution. For example, if your distribution is Oracle Linux 7 for x86_64, then select **Oracle Linux 7 Latest (x86_64)**.
7. Click **Subscribe**.
8. Start a terminal session and enter the following command as `root`, depending on your platform. For example:

- Oracle Linux 6 and Oracle Linux 7:

```
# yum install oracle-database-server-12cR2-preinstall
```

Note: Use the `-y` option if you want `yum` to skip the package confirmation prompt.

You should see output indicating that you have subscribed to the Oracle Linux channel, and that packages are being installed.

Oracle Linux automatically creates standard (not role-allocated) Oracle installation owner and groups, and sets up other kernel configuration settings as required for Oracle installations.

9. Enter the following command as `root` to update the `sysctl.conf` settings:

```
# sysctl -p
```

Note: The RPM packages set the Oracle software user to `oracle` by default. Before installing Oracle Database, you can update the Oracle user name in the `/etc/security/limits.d/oracle-database-server-12cR2-preinstall.conf` file and other configuration files.

10. Check the RPM log file to review the system configuration changes. For example, on Oracle Linux 7:

```
/var/log/oracle-database-server-12cR2-preinstall/backup/timestamp/orakernel.log
```

Repeat these steps on all other servers in your cluster.

3.4 Installing a New Oracle Linux Installation from DVDs or Images

Use this procedure to install a new Oracle Linux installation and to perform system configuration with the Oracle Preinstallation RPM:

1. Obtain Oracle Linux either by ordering the Oracle Linux media pack from Oracle Store, or by downloading disk images from the Oracle Software Delivery Cloud website for Oracle Linux and Oracle VM.
 - Oracle Store
<https://shop.oracle.com>
 - Oracle Software Delivery Cloud website:
<https://edelivery.oracle.com/linux>
2. Start the Oracle Linux installation and respond to installation screens with values appropriate for your environment.
3. Review the first software selection screen, which lists task-specific software options. At the bottom of the screen, there is an option to customize now or customize later. Select **Customize now**, and click **Next**.
4. On Oracle Linux, select **Servers** on the left of the screen and **System administration tools** on the right of the screen (options may vary between releases).

The Packages in System Tools window opens.

5. Select the Oracle Preinstallation RPM package box from the package list. For example, for Oracle Linux 7, select a package similar to the following:

```
oracle-database-server-12cR2-preinstall1-1.0-4.el7.x86_64.rpm
```

If you do not have an Oracle Preinstallation RPM package option that is current for your Oracle Database release, because you are using an Oracle Linux installation that is previous to your Oracle Database and Oracle Grid Infrastructure release, then install the RPM for your release manually after completing the operating system installation.

6. Close the optional package window and click **Next**.
7. Complete the other screens to finish the Oracle Linux installation.

Oracle Linux automatically creates a standard (not role-allocated) Oracle installation owner and groups, and sets up other kernel configuration settings as required for Oracle installations.

8. Repeat steps 2 through 6 on all other cluster member nodes.

3.5 Installing Oracle Linux with Oracle Linux Yum Server Support

Use this procedure to install Oracle Linux and configure your Linux installation for security errata or bug fix updates using the Oracle Linux yum server.

1. Obtain Oracle Linux DVDs from Oracle Store, or download Oracle Linux from the Oracle Software Delivery Cloud:

Oracle Store:

<https://shop.oracle.com>

Oracle Software Delivery Cloud website:

<https://edelivery.oracle.com/linux>

2. Install Oracle Linux from the ISO or DVD image.
3. Log in as `root`.
4. Download the yum repository file for your Linux distribution from the following URL using the instructions you can find on the Oracle Linux yum server website:

<http://yum.oracle.com/>

For example:

```
# cd /etc/yum.repos.d/
wget http://yum.oracle.com/public-yum-ol6.repo
```

Ensure that the `public_olrelease_latest` file (for example, `public_ol6_latest` for Oracle Linux 6) is enabled, because this is the repository that contains the Oracle Preinstallation RPM.

5. (Optional) Edit the `repo` file to enable other repositories. For example, enable the repository `public_ol6_latest` by setting `enabled=1` in the file with a text editor.
6. Run the command `yum repolist` to verify the registered channels.

7. Start a terminal session and enter the following command as `root`, depending on your platform. For example:

Oracle Linux 6 and Oracle Linux 7:

```
# yum install oracle-database-server-12cR2-preinstall
```

Note: Use the `-y` option if you want `yum` to skip the package confirmation prompt.

You should see output indicating that you have subscribed to the Oracle Linux channel, and that packages are being installed.

Oracle Linux automatically creates a standard (not role-allocated) Oracle installation owner and groups and sets up other kernel configuration settings as required for Oracle installations. If you plan to use job-role separation, then create the extended set of database users and groups depending on your requirements.

After installation, run the command `yum update` as needed to obtain the most current security errata and bug fixes for your Oracle Linux installation.

3.6 Configuring Oracle Ksplice to Perform Oracle Linux Updates

You can configure Oracle Ksplice to perform zero downtime updates for your operating system if you have an Oracle Linux Premier Support subscription and an access key, which is available on Unbreakable Linux Network (ULN).

For more information about Ksplice (including trial versions), refer to the Ksplice site:

<http://kssplice.oracle.com/>

To use Ksplice, you must register your system and create a Ksplice repository.

1. Check for your kernel distribution at the following URL:

<http://kssplice.oracle.com/legacy#supported-kernels>

2. Log in as `root`.

3. Ensure that you have access to the Internet on the server where you want to use Ksplice. For example, if you are using a proxy server, then set the proxy server and port values in the shell with commands similar to the following:

```
# export http_proxy=http://proxy.example.com:port
# export https_proxy=http://proxy.example.com:port
```

4. Download the Ksplice Uptrack repository RPM package:

<https://www.kssplice.com/yum/uptrack/ol/kssplice-uptrack-release.noarch.rpm>

5. Run the following commands:

```
rpm -i ksplice-uptrack-release.noarch.rpm
yum -y install uptrack
```

3.7 Configure Additional Operating System Features

Oracle recommends that you configure your operating system before starting installation with additional features, such as IPMI or additional programming environments.

Review the preinstallation chapters to ensure that you have completed configuration as needed for the features you want.

Configuring Operating Systems for Oracle Grid Infrastructure on Linux

Complete operating system configuration requirements and checks for Linux operating systems before you start installation.

[Guidelines for Linux Operating System Installation](#) (page 4-2)

Operating system guidelines to be aware of before proceeding with an Oracle installation.

[Reviewing Operating System and Software Upgrade Best Practices](#) (page 4-7)

These topics provide general planning guidelines and platform-specific information about upgrades and migration.

[Reviewing Operating System Security Common Practices](#) (page 4-9)

Secure operating systems are an important basis for general system security.

[About Installation Fixup Scripts](#) (page 4-9)

Oracle Universal Installer detects when the minimum requirements for an installation are not met, and creates shell scripts, called fixup scripts, to finish incomplete system configuration steps.

[About Operating System Requirements](#) (page 4-10)

Depending on the products that you intend to install, verify that you have the required operating system kernel and packages installed.

[Using Oracle RPM Checker on IBM: Linux on System z](#) (page 4-10)

Use the Oracle RPM Checker utility to verify that you have the required Red Hat Enterprise Linux or SUSE packages installed on the operating system before you start the Oracle Database or Oracle Grid Infrastructure installation.

[Operating System Requirements for x86-64 Linux Platforms](#) (page 4-11)

The Linux distributions and packages listed in this section are supported for this release on x86-64.

[Operating System Requirements for IBM: Linux on System z](#) (page 4-20)

The Linux distributions and packages listed in this section are supported for this release on IBM: Linux on System z.

[Additional Drivers and Software Packages for Linux](#) (page 4-23)

Information about optional drivers and software packages.

[Checking Kernel and Package Requirements for Linux](#) (page 4-27)

To check your kernel and packages to see if they meet minimum requirements for installation, perform the following steps:

[Installing the cvuqdisk RPM for Linux](#) (page 4-28)

If you do not use an Oracle Preinstallation RPM, and you want to use the Cluster Verification Utility, then you must install the `cvuqdisk` RPM.

[Reviewing HugePages Memory Allocation](#) (page 4-28)

Review this information if your operating system has HugePages enabled.

[Disabling Transparent HugePages](#) (page 4-29)

Oracle recommends that you disable Transparent HugePages before you start installation.

[Enabling the Name Service Cache Daemon](#) (page 4-30)

To allow Oracle Clusterware to better tolerate network failures with NAS devices or NFS mounts, enable the Name Service Cache Daemon (`nscd`).

[Verifying the Disk I/O Scheduler on Linux](#) (page 4-30)

For best performance for Oracle ASM, Oracle recommends that you use the Deadline I/O Scheduler.

[Using Automatic SSH Configuration During Installation](#) (page 4-31)

To install Oracle software, configure secure shell (SSH) connectivity between all cluster member nodes.

[Setting Network Time Protocol for Cluster Time Synchronization](#) (page 4-32)

Use either the Oracle Cluster Time Synchronization or network time protocol (NTP) option for time synchronization.

4.1 Guidelines for Linux Operating System Installation

Operating system guidelines to be aware of before proceeding with an Oracle installation.

This section provides information about installing a supported Linux distribution. Complete the minimum hardware configuration before you install the operating system.

[Completing a Minimal Linux Installation](#) (page 4-3)

Review information about minimal Linux operating system installation.

[Completing a Default Linux Installation](#) (page 4-4)

If you do not install the Oracle Preinstallation RPM, then Oracle recommends that you install your Linux operating system with the default software packages (RPMs).

[About Oracle Linux with the Unbreakable Enterprise Kernel](#) (page 4-4)

The Unbreakable Enterprise Kernel for Oracle Linux provides the latest innovations from upstream development to customers who run Oracle Linux in the data center.

[About the Oracle Preinstallation RPM](#) (page 4-5)

If your Linux distribution is Oracle Linux, or Red Hat Enterprise Linux, and you are an Oracle Linux customer, then you can complete most preinstallation configuration tasks by using the Oracle Preinstallation RPM for your release.

[Restrictions for HugePages and Transparent HugePages Configurations](#)
(page 4-6)

Review the following guidelines for HugePages and Transparent HugePages.

Related Topics:

[Oracle Linux Operating System Documentation](#)

4.1.1 Completing a Minimal Linux Installation

Review information about minimal Linux operating system installation.

Review the topics in this section if you intend to complete a minimal Linux installation.

[About Minimal Linux Installations](#) (page 4-3)

To complete a minimal Linux installation, select one of the minimal installation options.

[RPM Packages for Completing Operating System Configuration](#) (page 4-4)

You can use the following RPM packages to complete basic operating system configuration.

[Open SSH Requirement for Minimal Installation](#) (page 4-4)

OpenSSH is required for Oracle Grid Infrastructure installation.

4.1.1.1 About Minimal Linux Installations

To complete a minimal Linux installation, select one of the minimal installation options.

A minimal Linux installation option is either a custom installation where you select the **Minimal** option from Package Group Selection, or where you deselect all packages except for the Base pack. A minimal Linux installation lacks many RPMs required for database installation, so you must use an RPM package for your Oracle Linux release to install the required packages. The package you use depends on your Linux release, and your support status with Unbreakable Linux Network (ULN).

Note:

The Oracle Preinstallation RPM installs the X11 client libraries, but it does not install the X Window System server packages. To use graphical user interfaces such as OUI, configuration assistants, and Oracle Enterprise Manager, set the display to a system with X Window System server packages.

Refer to the following URL for documentation regarding installation of a reduced set of packages:

<https://support.oracle.com/rs?type=doc&id=579101.1>

If you are not a member of Unbreakable Linux Network or Red Hat Support network, and you are a My Oracle Support customer, then you can download instructions to configure a script that documents installation of a reduced set of packages from the same My Oracle Support Note.

You can also search for "Linux reduced set of packages" to locate the instructions.

4.1.1.2 RPM Packages for Completing Operating System Configuration

You can use the following RPM packages to complete basic operating system configuration.

Oracle Linux Preinstallation RPM with ULN Support

Oracle Preinstallation RPM for your Oracle Linux 6 or Oracle Linux 7 kernel (`oracle-rdbms-server-12cR2-preinstall`).

Unbreakable Linux Network (ULN) customers can obtain the Oracle Preinstallation RPM by using `yum`.

Oracle Linux Preinstallation RPM Without ULN Support

Oracle Linux packages on the Oracle Linux yum server website:

http://yum.oracle.com/repo/OracleLinux/OL6/latest/x86_64/

http://yum.oracle.com/repo/OracleLinux/OL7/latest/x86_64

4.1.1.3 Open SSH Requirement for Minimal Installation

OpenSSH is required for Oracle Grid Infrastructure installation.

OpenSSH should be included in the Linux distribution minimal installation. To confirm that SSH packages are installed, enter the following command:

```
# rpm -qa | grep ssh
```

If you do not see a list of SSH packages, then install those packages for your Linux distribution.

4.1.2 Completing a Default Linux Installation

If you do not install the Oracle Preinstallation RPM, then Oracle recommends that you install your Linux operating system with the default software packages (RPMs).

A default Linux installation includes most of the required packages and helps you limit manual verification of package dependencies. Oracle recommends that you do not customize the RPMs during installation.

For information about a default installation, log on to My Oracle Support:

<https://support.oracle.com>

Search for "default rpms linux installation," and look for your Linux distribution. For example:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=401167.1>

After installation, review system requirements for your distribution to ensure that you have all required kernel packages installed, and complete all other configuration tasks required for your distribution and system configuration.

4.1.3 About Oracle Linux with the Unbreakable Enterprise Kernel

The Unbreakable Enterprise Kernel for Oracle Linux provides the latest innovations from upstream development to customers who run Oracle Linux in the data center.

The Unbreakable Enterprise Kernel is included and enabled by default starting with Oracle Linux 5 Update 6 and later Oracle Linux kernels. It is based on a recent stable mainline development Linux kernel, and also includes optimizations developed in collaboration with Oracle Database, Oracle middleware, and Oracle hardware engineering teams to ensure stability and optimal performance for the most demanding enterprise workloads.

Oracle highly recommends deploying the Unbreakable Enterprise Kernel in your Oracle Linux environment, especially if you run enterprise applications. However, using Unbreakable Enterprise Kernel is optional. If you require strict Red Hat Enterprise Linux (RHEL) kernel compatibility, then Oracle Linux also includes a kernel compatible with the RHEL Linux kernel, compiled directly from the RHEL source code.

You can obtain more information about the Unbreakable Enterprise Kernel for Oracle Linux at the following URL:

<http://www.oracle.com/us/technologies/linux>

The Unbreakable Enterprise Kernel for Oracle Linux is the standard kernel used with Oracle products. The build and QA systems for Oracle Database and other Oracle products use the Unbreakable Enterprise Kernel for Oracle Linux exclusively. The Unbreakable Enterprise Kernel for Oracle Linux is also the kernel used in Oracle Exadata and Oracle Exalogic systems. Unbreakable Enterprise Kernel for Oracle Linux is used in all benchmark tests on Linux in which Oracle participates, as well as in the Oracle Preinstallation RPM program for x86-64.

Oracle Ksplice, which is part of Oracle Linux, updates the Linux operating system (OS) kernel, while it is running, without requiring restarts or any interruption. Ksplice is available only with Oracle Linux.

4.1.4 About the Oracle Preinstallation RPM

If your Linux distribution is Oracle Linux, or Red Hat Enterprise Linux, and you are an Oracle Linux customer, then you can complete most preinstallation configuration tasks by using the Oracle Preinstallation RPM for your release.

Oracle Preinstallation RPMs are available from the Oracle Linux Network or available on the Oracle Linux DVDs. Using the Oracle Preinstallation RPM is not required, but Oracle recommends you use it to save time in setting up your cluster servers.

When installed, the Oracle Preinstallation RPM does the following:

- Automatically downloads and installs any additional RPM packages needed for installing Oracle Grid Infrastructure and Oracle Database, and resolves any dependencies
- Creates an `oracle` user, and creates the `oraInventory` (`oinstall`) and `OSDBA` (`dba`) groups for that user
- As needed, sets `sysctl.conf` settings, system startup parameters, and driver parameters to values based on recommendations from the Oracle RDBMS Pre-Install program
- Sets hard and soft resource limits
- Sets other recommended parameters, depending on your kernel version

Note:

The Oracle Preinstallation RPM does not install OpenSSH, which is required for Oracle Grid Infrastructure installation. If you perform a minimal Linux installation and install the Oracle Preinstallation RPM for your release, then you must also install the OpenSSH client manually. Using RSH is no longer supported.

To become an Oracle Linux Network customer, contact your sales representative or purchase a license from the Oracle Linux store:

<https://shop.oracle.com/product/oraclelinux>

To register your server on the Unbreakable Linux Network, or to find out more information, see the following URL:

<https://linux.oracle.com>

If you are using Oracle Linux 5.2 or a higher release, then the Oracle Preinstallation RPM is included on the install media.

Note:

The Oracle Preinstallation RPM designated for each Oracle Database release sets kernel parameters and resource limits only for the user account `oracle`. To use multiple software account owners, you must perform system configuration for other accounts manually.

4.1.5 Restrictions for HugePages and Transparent HugePages Configurations

Review the following guidelines for HugePages and Transparent HugePages.

Oracle recommends that you disable Transparent HugePages, because they may cause delays in accessing memory that can result in node restarts in Oracle RAC environments, or performance issues or delays for Oracle Database single instances. Oracle continues to recommend using standard HugePages for Linux.

Transparent HugePages memory is enabled by default with Red Hat Enterprise Linux 6, Red Hat Enterprise Linux 7, SUSE 11, Oracle Linux 6, and Oracle Linux 7 with earlier releases of Oracle Linux with the Unbreakable Enterprise Kernel 2 (UEK2) kernels. Transparent HugePages memory is disabled by default in later releases of UEK2 and later UEK kernels.

Note: Although Transparent HugePages is disabled on UEK2 and later UEK kernels, Transparent HugePages may be enabled by default on your Linux system.

Transparent HugePages are similar to standard HugePages. However, while standard HugePages allocate memory at startup, Transparent HugePages memory uses the `khugepaged` thread in the kernel to allocate memory dynamically during runtime, using swappable HugePages.

HugePages allocates non-swappable memory for large page tables using memory-mapped files. HugePages are not enabled by default. If you enable HugePages, then you should deduct the memory allocated to HugePages from the available RAM

before calculating swap space. Refer to your distribution documentation and to Oracle Technology Network and My Oracle Support for more information.

During Oracle Grid Infrastructure installation, the Grid Infrastructure Management Repository (GIMR) is configured to use HugePages. Because the Grid Infrastructure Management Repository database starts before all other databases installed on the cluster, if the space allocated to HugePages is insufficient, then the System Global Area (SGA) of one or more databases may be mapped to regular pages, instead of Hugepages, which can adversely affect performance. Configure the HugePages memory allocation to a size large enough to accommodate the sum of the SGA sizes of all the databases you intend to install on the cluster, as well as the Grid Infrastructure Management Repository.

Related Topics:

Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems

4.2 Reviewing Operating System and Software Upgrade Best Practices

These topics provide general planning guidelines and platform-specific information about upgrades and migration.

[General Upgrade Best Practices](#) (page 4-7)

Be aware of these guidelines as a best practice before you perform an upgrade.

[New Server Operating System Upgrade Option](#) (page 4-8)

You can upgrade your operating system by installing a new operating system on a server, and then migrating your database using one of the following options:

[Oracle ASM Upgrade Notifications](#) (page 4-8)

Be aware of the following issues regarding Oracle ASM upgrades:

4.2.1 General Upgrade Best Practices

Be aware of these guidelines as a best practice before you perform an upgrade.

If you have an existing Oracle installation, then do the following:

- Record the version numbers, patches, and other configuration information
- Review upgrade procedures for your existing installation
- Review Oracle upgrade documentation before proceeding with installation, to decide how you want to proceed

Caution:

Always create a backup of existing databases before starting any configuration change.

Refer to *Oracle Database Upgrade Guide* for more information about required software updates, pre-upgrade tasks, post-upgrade tasks, compatibility, and interoperability between different releases.

Related Topics:

Oracle Database Upgrade Guide

4.2.2 New Server Operating System Upgrade Option

You can upgrade your operating system by installing a new operating system on a server, and then migrating your database using one of the following options:

Note:

Confirm that the server operating system is supported, and that kernel and package requirements for the operating system meet or exceed the minimum requirements for the Oracle Database release to which you want to migrate.

Manual, Command-Line Copy for Migrating Data and Upgrading Oracle Database

You can copy files to the new server and upgrade it manually. If you use this procedure, then you cannot use Oracle Database Upgrade Assistant. However, you can revert to your existing database if you encounter upgrade issues.

1. Copy the database files from the computer running the previous operating system to the one running the new operating system.
2. Re-create the control files on the computer running the new operating system.
3. Manually upgrade the database using command-line scripts and utilities.

See Also:

Oracle Database Upgrade Guide to review the procedure for upgrading the database manually, and to evaluate the risks and benefits of this option

Export/Import Method for Migrating Data and Upgrading Oracle Database

You can install the operating system on the new server, install the new Oracle Database release on the new server, and then use Oracle Data Pump Export and Import utilities to migrate a copy of data from your current database to a new database in the new release. Data Pump Export and Import are recommended for higher performance and to ensure support for new data types.

See Also:

Oracle Database Upgrade Guide to review the Export/Import method for migrating data and upgrading Oracle Database

4.2.3 Oracle ASM Upgrade Notifications

Be aware of the following issues regarding Oracle ASM upgrades:

- You can upgrade Oracle Automatic Storage Management (Oracle ASM) 11g release 2 (11.2) and later without shutting down an Oracle RAC database by performing a rolling upgrade either of individual nodes, or of a set of nodes in the

cluster. However, if you have a standalone database on a cluster that uses Oracle ASM, then you must shut down the standalone database before upgrading.

- The location of the Oracle ASM home changed in Oracle Grid Infrastructure 11g release 2 (11.2) so that Oracle ASM is installed with Oracle Clusterware in the Oracle Grid Infrastructure home (Grid home).
- Two nodes of different releases cannot run in the cluster. When upgrading from Oracle Grid Infrastructure 11g release 2 (11.2) or Oracle Grid Infrastructure 12c release 1 (12.1) to a later release, if there is an outage during the rolling upgrade, then when you restart the upgrade, ensure that you start the earlier release of Oracle Grid Infrastructure and bring the Oracle ASM cluster back in the rolling migration mode.

4.3 Reviewing Operating System Security Common Practices

Secure operating systems are an important basis for general system security.

Ensure that your operating system deployment is in compliance with common security practices as described in your operating system vendor security guide.

4.4 About Installation Fixup Scripts

Oracle Universal Installer detects when the minimum requirements for an installation are not met, and creates shell scripts, called fixup scripts, to finish incomplete system configuration steps.

If Oracle Universal Installer detects an incomplete task, then it generates fixup scripts (`runfixup.sh`). You can run the fixup script and click **Fix and Check Again**. The fixup script modifies both persistent parameter settings and parameters in memory, so you do not have to restart the system.

The Fixup script does the following tasks:

- Sets kernel parameters, if necessary, to values required for successful installation, including:
 - Shared memory parameters.
 - Open file descriptor and UDP send/receive parameters.
- Creates and sets permissions on the Oracle Inventory (central inventory) directory.
- Creates or reconfigures primary and secondary group memberships for the installation owner, if necessary, for the Oracle Inventory directory and the operating system privileges groups.
- Sets shell limits, if necessary, to required values.

Note:

Using fixup scripts does not ensure that all the prerequisites for installing Oracle Database are met. You must still verify that all the preinstallation requirements are met to ensure a successful installation.

Oracle Universal Installer is fully integrated with Cluster Verification Utility (CVU) automating many prerequisite checks for your Oracle Grid Infrastructure or Oracle

Real Application Clusters (Oracle RAC) installation. You can also manually perform various CVU verifications by running the `cluvfy` command.

Related Topics:

Oracle Clusterware Administration and Deployment Guide

4.5 About Operating System Requirements

Depending on the products that you intend to install, verify that you have the required operating system kernel and packages installed.

Requirements listed in this document are current as of the date listed on the title page. To obtain the most current information about kernel requirements, see the online version at the following URL:

<http://docs.oracle.com>

Oracle Universal Installer performs checks on your system to verify that it meets the listed operating system package requirements. To ensure that these checks complete successfully, verify the requirements before you start OUI.

Note:

Oracle does not support running different operating system versions on cluster members, unless an operating system is being upgraded. You cannot run different operating system version binaries on members of the same cluster, even if each operating system is supported.

4.6 Using Oracle RPM Checker on IBM: Linux on System z

Use the Oracle RPM Checker utility to verify that you have the required Red Hat Enterprise Linux or SUSE packages installed on the operating system before you start the Oracle Database or Oracle Grid Infrastructure installation.

Download the Oracle RPM Checker utility from the link in the My Oracle Support note 1574412.1:

<https://support.oracle.com/rs?type=doc&id=1574412.1>

Download the Oracle RPM Checker utility for your IBM: Linux on System z distribution, unzip the RPM, and install the RPM as `root`. Then, run the utility as `root` to check your operating system packages. For example:

On Red Hat Enterprise Linux 6:

```
# rpm -ivh ora-val-rpm-EL6-DB-12.1.0.1-1.s390x.rpm
```

On Red Hat Enterprise Linux 7:

```
# rpm -ivh ora-val-rpm-RH7-DB-12.1.0.2-1.s390x.rpm
```

On Red Hat Enterprise Linux, the utility checks and also installs all required RPMs. For example:

On Red Hat Enterprise Linux 6:

```
# yum install ora-val-rpm-EL6-DB-12.1.0.1-1.s390x.rpm
```

On Red Hat Enterprise Linux 7:

```
# yum install ora-val-rpm-RH7-DB-12.1.0.2-1.s390x.rpm
```

4.7 Operating System Requirements for x86-64 Linux Platforms

The Linux distributions and packages listed in this section are supported for this release on x86-64.

Identify the requirements for your Linux distribution, and ensure that you have a supported kernel and required packages installed before starting installation.

Note:

- Oracle Universal Installer requires an X Window System (for example, `libx`). The `libx` packages are part of a default Linux installation. If you perform an install on a system with a reduced set of packages, then you must ensure that `libx` or a similar X Window System package is installed.
 - The Unbreakable Enterprise Kernel for Oracle Linux can be installed on x86-64 servers running either Oracle Linux or Red Hat Enterprise Linux. As of Oracle Linux 5 Update 6, the Unbreakable Enterprise Kernel is the default system kernel. An x86 (32-bit) release of Oracle Linux including the Unbreakable Enterprise Kernel is available with Oracle Linux 5 update 7 and later.
 - 32-bit packages in these requirements lists are needed only if you intend to use 32-bit client applications to access 64-bit servers.
 - Oracle Database 12c Release 2 (12.2) and later does not require the compiler packages `gcc` and `gcc-c++` on Oracle Linux and Red Hat Enterprise Linux for Oracle Database or Oracle Grid Infrastructure installations.
-
-

The platform-specific hardware and software requirements included in this guide were current when this guide was published. However, because new platforms and operating system software versions may be certified after this guide is published, review the certification matrix on the My Oracle Support website for the most up-to-date list of certified hardware platforms and operating system versions:

<https://support.oracle.com/>

[Supported Oracle Linux 7 Distributions for x86-64](#) (page 4-12)

Use the following information to check supported Oracle Linux 7 distributions:

[Supported Oracle Linux 6 Distributions for x86-64](#) (page 4-13)

Use the following information to check supported Oracle Linux 6 distributions:

[Supported Red Hat Enterprise Linux 7 Distributions for x86-64](#) (page 4-15)

Use the following information to check supported Red Hat Enterprise Linux 7 distributions:

[Supported Red Hat Enterprise Linux 6 Distributions for x86-64](#) (page 4-16)

Use the following information to check supported Red Hat Enterprise Linux 6 distributions:

[Supported SUSE Linux Enterprise Server 12 Distributions for x86-64](#) (page 4-17)
 Use the following information to check supported SUSE Linux Enterprise Server 12 distributions:

See Also:

If you currently use, or plan to upgrade to, Oracle Linux 7.2 or Red Hat Enterprise Linux 7.2, then see information about the RemoveIPC settings:

- My Oracle Support Note 2081410.1:
<https://support.oracle.com/rs?type=doc&id=2081410.1>
 - Oracle Linux 7 Update 2 Release Notes:
<http://docs.oracle.com/en/operating-systems/>
-
-

4.7.1 Supported Oracle Linux 7 Distributions for x86-64

Use the following information to check supported Oracle Linux 7 distributions:

Table 4-1 x86-64 Oracle Linux 7 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Oracle Linux 7	Subscribe to the Oracle Linux 7 channel on the Unbreakable Linux Network, or configure a yum repository from the Oracle Linux yum server website, and then install the Oracle Preinstallation RPM. This RPM installs all required kernel packages for Oracle Grid Infrastructure and Oracle Database installations, and performs other system configuration. Supported distributions: <ul style="list-style-type: none"> • Oracle Linux 7 with the Unbreakable Enterprise Kernel 3: 3.8.13-35.3.1.el7uek.x86_64 or later • Oracle Linux 7.2 with the Unbreakable Enterprise Kernel 4: 4.1.12-32.2.3.el7uek.x86_64 or later • Oracle Linux 7 with the Red Hat Compatible kernel: 3.10.0-123.el7.x86_64 or later

Table 4-1 (Cont.) x86-64 Oracle Linux 7 Minimum Operating System Requirements

Item	Requirements
Packages for Oracle Linux 7	<p>The following packages (or later versions) must be installed:</p> <p>binutils-2.23.52.0.1-12.el7 (x86_64) compat-libcap1-1.10-3.el7 (x86_64) compat-libstdc++-33-3.2.3-71.el7 (i686) compat-libstdc++-33-3.2.3-71.el7 (x86_64) glibc-2.17-36.el7 (i686) glibc-2.17-36.el7 (x86_64) glibc-devel-2.17-36.el7 (i686) glibc-devel-2.17-36.el7 (x86_64) ksh libaio-0.3.109-9.el7 (i686) libaio-0.3.109-9.el7 (x86_64) libaio-devel-0.3.109-9.el7 (i686) libaio-devel-0.3.109-9.el7 (x86_64) libX11-1.6.0-2.1.el7 (i686) libX11-1.6.0-2.1.el7 (x86_64) libXau-1.0.8-2.1.el7 (i686) libXau-1.0.8-2.1.el7 (x86_64) libXi-1.7.2-1.el7 (i686) libXi-1.7.2-1.el7 (x86_64) libXtst-1.2.2-1.el7 (i686) libXtst-1.2.2-1.el7 (x86_64) libgcc-4.8.2-3.el7 (i686) libgcc-4.8.2-3.el7 (x86_64) libstdc++-4.8.2-3.el7 (i686) libstdc++-4.8.2-3.el7 (x86_64) libstdc++-devel-4.8.2-3.el7 (i686) libstdc++-devel-4.8.2-3.el7 (x86_64) libxcb-1.9-5.el7 (i686) libxcb-1.9-5.el7 (x86_64) make-3.82-19.el7 (x86_64) nfs-utils-1.3.0-0.21.el7.x86_64 (for Oracle ACFS) net-tools-2.0-0.17.20131004git.el7 (x86_64) (for Oracle RAC and Oracle Clusterware) smartmontools-6.2-4.el7 (x86_64) sysstat-10.1.5-1.el7 (x86_64)</p>

4.7.2 Supported Oracle Linux 6 Distributions for x86-64

Use the following information to check supported Oracle Linux 6 distributions:

Table 4-2 x86-64 Oracle Linux 6 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.

Table 4-2 (Cont.) x86-64 Oracle Linux 6 Minimum Operating System Requirements

Item	Requirements
Oracle Linux 6	<p data-bbox="626 323 1357 495">Subscribe to the Oracle Linux 6 channel on the Unbreakable Linux Network, or configure a yum repository from the Oracle Linux yum server website, and then install the Oracle Preinstallation RPM. This RPM installs all required kernel packages for Oracle Grid Infrastructure and Oracle Database installations, and performs other system configuration.</p> <p data-bbox="626 506 894 531">Supported distributions:</p> <ul data-bbox="626 541 1312 779" style="list-style-type: none"><li data-bbox="626 541 1312 596">• Oracle Linux 6.4 with the Unbreakable Enterprise Kernel 2: 2.6.39-400.211.1.el6uek.x86_64 or later<li data-bbox="626 606 1312 661">• Oracle Linux 6.6 with the Unbreakable Enterprise Kernel 3: 3.8.13-44.1.1.el6uek.x86_64 or later<li data-bbox="626 672 1312 726">• Oracle Linux 6.8 with the Unbreakable Enterprise Kernel 4: 4.1.12-37.6.2.el6uek.x86_64 or later<li data-bbox="626 737 1312 779">• Oracle Linux 6.4 with the Red Hat Compatible kernel: 2.6.32-358.el6.x86_64 or later

Table 4-2 (Cont.) x86-64 Oracle Linux 6 Minimum Operating System Requirements

Item	Requirements
Packages for Oracle Linux 6	<p>The following packages (or later versions) must be installed:</p> <ul style="list-style-type: none"> binutils-2.20.51.0.2-5.36.el6 (x86_64) compat-libcap1-1.10-1 (x86_64) compat-libstdc++-33-3.2.3-69.el6 (x86_64) compat-libstdc++-33-3.2.3-69.el6 (i686) e2fsprogs-1.41.12-14.el6 (x86_64) e2fsprogs-libs-1.41.12-14.el6 (x86_64) glibc-2.12-1.7.el6 (i686) glibc-2.12-1.7.el6 (x86_64) glibc-devel-2.12-1.7.el6 (x86_64) glibc-devel-2.12-1.7.el6 (i686) ksh libgcc-4.4.4-13.el6 (i686) libgcc-4.4.4-13.el6 (x86_64) libstdc++-4.4.4-13.el6 (x86_64) libstdc++-4.4.4-13.el6 (i686) libstdc++-devel-4.4.4-13.el6 (x86_64) libstdc++-devel-4.4.4-13.el6 (i686) libaio-0.3.107-10.el6 (x86_64) libaio-0.3.107-10.el6 (i686) libaio-devel-0.3.107-10.el6 (x86_64) libaio-devel-0.3.107-10.el6 (i686) libXtst-1.0.99.2 (x86_64) libXtst-1.0.99.2 (i686) libX11-1.5.0-4.el6 (i686) libX11-1.5.0-4.el6 (x86_64) libXau-1.0.6-4.el6 (i686) libXau-1.0.6-4.el6 (x86_64) libxcb-1.8.1-1.el6 (i686) libxcb-1.8.1-1.el6 (x86_64) libXi-1.3 (x86_64) libXi-1.3 (i686) make-3.81-19.el6 net-tools-1.60-110.el6_2.x86_64 (for Oracle RAC and Oracle Clusterware) nfs-utils-1.2.3-15.0.1 (for Oracle ACFS) sysstat-9.0.4-11.el6 (x86_64) smartmontools-5.43-1.el6.x86_64

4.7.3 Supported Red Hat Enterprise Linux 7 Distributions for x86-64

Use the following information to check supported Red Hat Enterprise Linux 7 distributions:

Table 4-3 x86-64 Red Hat Enterprise Linux 7 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Red Hat Enterprise Linux 7	Supported distributions: <ul style="list-style-type: none"> Red Hat Enterprise Linux 7: 3.10.0-123.el7.x86_64 or later
Packages for Red Hat Enterprise Linux 7	The following packages (or later versions) must be installed: <ul style="list-style-type: none"> binutils-2.23.52.0.1-12.el7 (x86_64) compat-libcap1-1.10-3.el7 (x86_64) compat-libstdc++-33-3.2.3-71.el7 (i686) compat-libstdc++-33-3.2.3-71.el7 (x86_64) glibc-2.17-36.el7 (i686) glibc-2.17-36.el7 (x86_64) glibc-devel-2.17-36.el7 (i686) glibc-devel-2.17-36.el7 (x86_64) ksh libaio-0.3.109-9.el7 (i686) libaio-0.3.109-9.el7 (x86_64) libaio-devel-0.3.109-9.el7 (i686) libaio-devel-0.3.109-9.el7 (x86_64) libgcc-4.8.2-3.el7 (i686) libgcc-4.8.2-3.el7 (x86_64) libstdc++-4.8.2-3.el7 (i686) libstdc++-4.8.2-3.el7 (x86_64) libstdc++-devel-4.8.2-3.el7 (i686) libstdc++-devel-4.8.2-3.el7 (x86_64) libxcb-1.9-5.el7 (i686) libxcb-1.9-5.el7 (x86_64) libX11-1.6.0-2.1.el7 (i686) libX11-1.6.0-2.1.el7 (x86_64) libXau-1.0.8-2.1.el7 (i686) libXau-1.0.8-2.1.el7 (x86_64) libXi-1.7.2-1.el7 (i686) libXi-1.7.2-1.el7 (x86_64) libXtst-1.2.2-1.el7 (i686) libXtst-1.2.2-1.el7 (x86_64) make-3.82-19.el7 (x86_64) net-tools-2.0-0.17.20131004git.el7 (x86_64) (for Oracle RAC and Oracle Clusterware) nfs-utils-1.3.0-0.21.el7.x86_64 (for Oracle ACFS) smartmontools-6.2-4.el7 (x86_64) sysstat-10.1.5-1.el7 (x86_64)

4.7.4 Supported Red Hat Enterprise Linux 6 Distributions for x86-64

Use the following information to check supported Red Hat Enterprise Linux 6 distributions:

Table 4-4 x86-64 Red Hat Enterprise Linux 6 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Red Hat Enterprise Linux 6	Supported distributions: <ul style="list-style-type: none"> Red Hat Enterprise Linux 6.4: 2.6.32-358.el6.x86_64 or later
Packages for Red Hat Enterprise Linux 6	The following packages (or later versions) must be installed: <ul style="list-style-type: none"> binutils-2.20.51.0.2-5.36.el6 (x86_64) compat-libcap1-1.10-1 (x86_64) compat-libstdc++-33-3.2.3-69.el6 (x86_64) compat-libstdc++-33-3.2.3-69.el6 (i686) e2fsprogs-1.41.12-14.el6 (x86_64) e2fsprogs-libs-1.41.12-14.el6 (x86_64) glibc-2.12-1.107.el6 (i686) glibc-2.12-1.107.el6 (x86_64) glibc-devel-2.12-1.107.el6 (i686) glibc-devel-2.12-1.107.el6 (x86_64) ksh libaio-0.3.107-10.el6 (x86_64) libaio-0.3.107-10.el6 (i686) libaio-devel-0.3.107-10.el6 (x86_64) libaio-devel-0.3.107-10.el6 (i686) libX11-1.5.0-4.el6 (i686) libX11-1.5.0-4.el6 (x86_64) libXau-1.0.6-4.el6 (i686) libXau-1.0.6-4.el6 (x86_64) libXi-1.6.1-3.el6 (i686) libXi-1.6.1-3.el6 (x86_64) libXtst-1.2.1-2.el6 (i686) libXtst-1.2.1-2.el6 (x86_64) libgcc-4.4.7-3.el6 (i686) libgcc-4.4.7-3.el6 (x86_64) libstdc++-4.4.7-3.el6 (i686) libstdc++-4.4.7-3.el6 (x86_64) libstdc++-devel-4.4.7-3.el6 (i686) libstdc++-devel-4.4.7-3.el6 (x86_64) libxcb-1.8.1-1.el6 (i686) libxcb-1.8.1-1.el6 (x86_64) make-3.81-20.el6 (x86_64) net-tools-1.60-110.el6_2 (x86_64) (for Oracle RAC and Oracle Clusterware) nfs-utils-1.2.3-36.el6 (x86_64) (for Oracle ACFS) smartmontools-5.43-1.el6 (x86_64) sysstat-9.0.4-20.el6 (x86_64)

4.7.5 Supported SUSE Linux Enterprise Server 12 Distributions for x86-64

Use the following information to check supported SUSE Linux Enterprise Server 12 distributions:

Table 4-5 x86-64 SUSE Linux Enterprise Server 12 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
SUSE Linux Enterprise Server	Supported distributions: SUSE Linux Enterprise Server 12 SP1: 3.12.49-11.1 or later

Table 4-5 (Cont.) x86-64 SUSE Linux Enterprise Server 12 Minimum Operating System Requirements

Item	Requirements
Package requirements for SUSE Linux Enterprise Server 12	<p>The following packages (or later versions) must be installed:</p> <ul style="list-style-type: none"> binutils-2.24-2.165.x86_64 gcc-c++-32bit-4.8-6.189.x86_64 gcc-c++-4.8-6.189.x86_64 gcc48-c++-4.8.3+r212056-6.3.x86_64 gcc-32bit-4.8-6.189.x86_64 gcc-4.8-6.189.x86_64 gcc-info-4.8-6.189.x86_64 gcc-locale-4.8-6.189.x86_64 gcc48-32bit-4.8.3+r212056-6.3.x86_64 gcc48-4.8.3+r212056-6.3.x86_64 gcc48-info-4.8.3+r212056-6.3.noarch gcc48-locale-4.8.3+r212056-6.3.x86_64 glibc-2.19-17.72.x86_64 glibc-devel-2.19-17.72.x86_64 libaio-devel-0.3.109-17.15.x86_64 libaio1-0.3.109-17.15.x86_64 libaio1-32bit-0.3.109-17.15.x86_64 libgfortran3-4.8.3+r212056-6.3.x86_64 libX11-6-1.6.2-4.12.x86_64 libX11-6-32bit-1.6.2-4.12.x86_64 libXau6-1.0.8-4.58.x86_64 libXau6-32bit-1.0.8-4.58.x86_64 libXtst6-1.2.2-3.60.x86_64 libXtst6-32bit-1.2.1-2.4.1.x86_64 libcap-ng-utils-0.7.3-4.125.x86_64 libcap-ng0-0.7.3-4.125.x86_64 libcap-ng0-32bit-0.7.3-4.125.x86_64 libcap-progs-2.22-11.709.x86_64 libcap1-1.10-59.61.x86_64 libcap1-32bit-1.10-59.61.x86_64 libcap2-2.22-11.709.x86_64 libcap2-32bit-2.22-11.709.x86_64 libgcc_s1-32bit-4.8.3+r212056-6.3.x86_64 libgcc_s1-4.8.3+r212056-6.3.x86_64 libpcap1-1.5.3-2.18.x86_64 libstdc++6-32bit-4.8.3+r212056-6.3.x86_64 libstdc++6-4.8.3+r212056-6.3.x86_64 make-4.0-2.107.x86_64 mksh-50-2.13.x86_64 net-tools-1.60-764.185.x86_64 (for Oracle RAC and Oracle Clusterware) nfs-kernel-server-1.3.0-6.9.x86_64 (for Oracle ACFS) smartmontools-6.2-4.33.x86_64 sysstat-8.1.5-7.32.1.x86_64 xorg-x11-libs-7.6-45.14

4.8 Operating System Requirements for IBM: Linux on System z

The Linux distributions and packages listed in this section are supported for this release on IBM: Linux on System z.

Identify the requirements for your IBM: Linux on System z distribution, and ensure that you have a supported kernel and required packages installed before starting installation.

Note: 32-bit packages in these requirements lists are needed only if you intend to use 32-bit client applications to access 64-bit servers.

The platform-specific hardware and software requirements included in this guide were current when this guide was published. However, because new platforms and operating system software versions may be certified after this guide is published, review the certification matrix on the My Oracle Support website for the most up-to-date list of certified hardware platforms and operating system versions:

<https://support.oracle.com/>

[Supported Red Hat Enterprise Linux 7 Distributions for IBM: Linux on System z](#)
(page 4-20)

Use the following information to check supported Red Hat Enterprise Linux 7 distributions:

[Supported Red Hat Enterprise Linux 6 Distributions for IBM: Linux on System z](#)
(page 4-21)

Use the following information to check supported Red Hat Enterprise Linux 6 distributions:

[Supported SUSE Linux Enterprise Server 12 Distributions for IBM: Linux on System z](#) (page 4-22)

Use the following information to check supported SUSE Linux Enterprise Server 12 distributions:

4.8.1 Supported Red Hat Enterprise Linux 7 Distributions for IBM: Linux on System z

Use the following information to check supported Red Hat Enterprise Linux 7 distributions:

Table 4-6 Red Hat Enterprise Linux 7 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Red Hat Enterprise Linux 7	Red Hat Enterprise Linux 7.2: 3.10.0-327.el7.s390x or later

Table 4-6 (Cont.) Red Hat Enterprise Linux 7 Minimum Operating System Requirements

Item	Requirements
Packages for Red Hat Enterprise Linux 7	The following packages (or later versions) must be installed: binutils-2.23.52.0.1-55.el7 (s390x) compat-libcap1-1.10-7.el7 (s390x) gcc-4.8.5-4.el7 (s390x) gcc-c++-4.8.5-4.el7 (s390x) glibc-2.17-105.el7 (s390) glibc-2.17-105.el7 (s390x) glibc-devel-2.17-105.el7 (s390) glibc-devel-2.17-105.el7 (s390x) ksh-20120801-22.el7_1.2 (s390x) libXaw-1.0.12-5.el7 (s390) libXft-2.3.2-2.el7 (s390) libXi-1.7.4-2.el7 (s390) libXi-1.7.4-2.el7 (s390x) libXmu-1.1.2-2.el7 (s390) libXp-1.0.2-2.1.el7 (s390) libXtst-1.2.2-2.1.el7 (s390) libXtst-1.2.2-2.1.el7 (s390x) libaio-0.3.109-13.el7 (s390) libaio-0.3.109-13.el7 (s390x) libaio-devel-0.3.109-13.el7 (s390x) libgcc-4.8.5-4.el7 (s390) libgcc-4.8.5-4.el7 (s390x) libgfortran-4.8.5-4.el7 (s390x) libstdc++-4.8.5-4.el7 (s390x) libstdc++-devel-4.8.5-4.el7 (s390x) make-3.82-21.el7 (s390x) sysstat-10.1.5-7.el7 (s390x)

4.8.2 Supported Red Hat Enterprise Linux 6 Distributions for IBM: Linux on System z

Use the following information to check supported Red Hat Enterprise Linux 6 distributions:

Table 4-7 Red Hat Enterprise Linux 6 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Red Hat Enterprise Linux 6	Red Hat Enterprise Linux 6.6: 2.6.32-504.el6.s390x or later

Table 4-7 (Cont.) Red Hat Enterprise Linux 6 Minimum Operating System Requirements

Item	Requirements
Packages for Red Hat Enterprise Linux 6	<p>The following packages (or later versions) must be installed:</p> <p>binutils-2.20.51.0.2-5.42.el6 (s390x) compat-libcap1-1.10-1 (s390x) compat-libstdc++-33-3.2.3-69.el6 (s390) compat-libstdc++-33-3.2.3-69.el6 (s390x) gcc-4.4.7-11.el6 (s390x) gcc-c++-4.4.7-11.el6 (s390x) glibc-static-2.12-1.149.el6 (s390x) glibc-2.12-1.149.el6 (s390) glibc-2.12-1.149.el6 (s390x) glibc-devel-2.12-1.149.el6 (s390) glibc-devel-2.12-1.149.el6 (s390x) ksh-20120801-21.el6 (s390x) libXaw-1.0.11-2.el6 (s390) libXft-2.3.1-2.el6 (s390) libXi-1.7.2-2.2.el6 (s390) libXi-1.7.2-2.2.el6 (s390x) libXmu-1.1.1-2.el6 (s390) libXp-1.0.2-2.1.el6 (s390) libXtst-1.2.2-2.1.el6 (s390) libXtst-1.2.2-2.1.el6 (s390x) libaio-0.3.107-10.el6 (s390) libaio-0.3.107-10.el6 (s390x) libaio-devel-0.3.107-10.el6 (s390x) libgcc-4.4.7-11.el6 (s390) libgcc-4.4.7-11.el6 (s390x) libstdc++-4.4.7-11.el6 (s390) libstdc++-4.4.7-11.el6 (s390x) libstdc++-devel-4.4.7-11.el6 (s390x) make-3.81-20.el6 (s390x) sysstat-9.0.4-27.el6 (s390x)</p>

4.8.3 Supported SUSE Linux Enterprise Server 12 Distributions for IBM: Linux on System z

Use the following information to check supported SUSE Linux Enterprise Server 12 distributions:

Table 4-8 SUSE Linux Enterprise Server 12 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
SUSE Linux Enterprise Server 12	SUSE Linux Enterprise Server 12 SP1: 3.12.49-11-default s390x or later

Table 4-8 (Cont.) SUSE Linux Enterprise Server 12 Minimum Operating System Requirements

Item	Requirements
Packages for SUSE Linux Enterprise Server 12	<p>The following packages (or later versions) must be installed:</p> <p>binutils-2.25.0-13.1 (s390x) gcc-4.8-6.189 (s390x) gcc-c++-4.8-6.189 (s390x) glibc-2.19-31.9 (s390x) glibc-32bit-2.19-31.9 (s390x) glibc-devel-2.19-31.9 (s390x) glibc-devel-32bit-2.19-31.9 (s390x) libaio-devel-0.3.109-17.15 (s390x) libaio-devel-32bit-0.3.109-17.15 libaio1-0.3.109-17.15 (s390x) libaio1-32bit-0.3.109-17.15 (s390x) (DOUBT) libX11-6-1.6.2-4.12 (s390x) libX11-6-32bit-1.6.2-4.12 (s390x) libXau6-1.0.8-4.58 (s390x) libXau6-32bit-1.0.8-4.58 (s390x) libXaw7-1.0.12-3.62 (s390x) libXext6-1.3.2-3.61 (s390x) libXext6-32bit-1.3.2-3.61 (s390x) libXft2-2.3.1-9.32 (s390x) libXft2-32bit-2.3.1-9.32 (s390x) libXi6-1.7.4-9.2.s390x libXi6-1.7.4-9.2 (s390x) libXi6-32bit-1.7.4-9.2 (s390x) libXmu6-1.1.2-3.60 (s390x) libXp6-1.0.2-3.58 (s390x) libXp6-1.0.2-3.58.s390x (32bit) libXtst6-1.2.2-3.60 (s390x) libXtst6-32bit-1.2.2-3.60 (s390x) libXmu6-1.1.2-3.60 (s390x) (32bit) libcap2-2.22-11.709 (s390x) libstdc++48-devel-32bit-4.8.5-24.1 (s390x) (DOUBT) libstdc++48-devel-4.8.5-24.1 (s390x) libstdc++6-32bit-5.2.1+r226025-4.1 (s390x) libstdc++6-5.2.1+r226025-4.1 (s390x) libxcb1-1.10-1.21 (s390x) libxcb1-32bit-1.10-1.21 (s390x) make-4.0-4.1 (s390x) mksh-50-2.13 (s390x)</p>

4.9 Additional Drivers and Software Packages for Linux

Information about optional drivers and software packages.

You are not required to install additional drivers and packages, but you may choose to install or configure these drivers and packages.

[Installing PAM for Login Authentication on Linux](#) (page 4-24)

Pluggable Authentication Modules (PAM) is a system of libraries that handle user authentication tasks for applications.

[Installation Requirements for OCFS2](#) (page 4-24)

Review the following sections to install OCFS2.

[Installation Requirements for ODBC and LDAP](#) (page 4-25)

Review these topics to install Open Database Connectivity (ODBC) and Lightweight Directory Access Protocol (LDAP).

[Installing Oracle Messaging Gateway](#) (page 4-26)

Oracle Messaging Gateway is installed with Enterprise Edition of Oracle Database. However, you may require a CSD or Fix Packs.

[Installation Requirements for Programming Environments for Linux](#) (page 4-26)

Review the following section to install programming environments:

[Installation Requirements for Web Browsers](#) (page 4-27)

Web browsers are required only if you intend to use Oracle Enterprise Manager Database Express and Oracle Enterprise Manager Cloud Control. Web browsers must support JavaScript, and the HTML 4.0 and CSS 1.0 standards.

4.9.1 Installing PAM for Login Authentication on Linux

Pluggable Authentication Modules (PAM) is a system of libraries that handle user authentication tasks for applications.

On Linux, external scheduler jobs require PAM. Oracle strongly recommends that you install the latest Linux-PAM library for your Linux distribution.

Use a package management system (`yum`, `up2date`, `YaST`) for your distribution to install the latest `pam` (Pluggable Authentication Modules for Linux) library.

4.9.2 Installation Requirements for OCFS2

Review the following sections to install OCFS2.

[About OCFS2 and Shared Storage](#) (page 4-24)

You can use Oracle Cluster File System 2 (OCFS2) with Oracle Grid Infrastructure. However, you are not required to use OCFS2.

[Installing OCFS2](#) (page 4-24)

OCFS2 Release 1.6 is included with the Unbreakable Enterprise Kernel available with Oracle Linux 6 and Oracle Linux 7.

4.9.2.1 About OCFS2 and Shared Storage

You can use Oracle Cluster File System 2 (OCFS2) with Oracle Grid Infrastructure. However, you are not required to use OCFS2.

OCFS2 is a POSIX-compliant general purpose shared disk cluster file system for Linux. OCFS2 is supported for this release only with Oracle Linux 7 and Oracle Linux 6.

4.9.2.2 Installing OCFS2

OCFS2 Release 1.6 is included with the Unbreakable Enterprise Kernel available with Oracle Linux 6 and Oracle Linux 7.

See the OCFS2 project page for additional information:

<http://oss.oracle.com/projects/ocfs2/>

4.9.3 Installation Requirements for ODBC and LDAP

Review these topics to install Open Database Connectivity (ODBC) and Lightweight Directory Access Protocol (LDAP).

[About ODBC Drivers and Oracle Database](#) (page 4-25)

Open Database Connectivity (ODBC) is a set of database access APIs that connect to the database, prepare, and then run SQL statements on the database.

[Installing ODBC Drivers for Linux x86-64](#) (page 4-25)

If you intend to use ODBC, then install the most recent ODBC Driver Manager for Linux.

[About LDAP and Oracle Plug-ins](#) (page 4-25)

Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over IP networks.

[Installing the LDAP Package](#) (page 4-25)

LDAP is included in a default Linux operating system installation.

4.9.3.1 About ODBC Drivers and Oracle Database

Open Database Connectivity (ODBC) is a set of database access APIs that connect to the database, prepare, and then run SQL statements on the database.

An application that uses an ODBC driver can access non-uniform data sources, such as spreadsheets and comma-delimited files.

4.9.3.2 Installing ODBC Drivers for Linux x86-64

If you intend to use ODBC, then install the most recent ODBC Driver Manager for Linux.

Download and install the ODBC Driver Manager and Linux RPMs from the following website:

<http://www.unixodbc.org>

Review the minimum supported ODBC driver releases, and install ODBC drivers of the following or later releases for all Linux distributions:

unixODBC-2.3.1 or later

4.9.3.3 About LDAP and Oracle Plug-ins

Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over IP networks.

You require the LDAP package if you want to use features requiring LDAP, including the Oracle Database scripts `odisrvreg` and `oidca` for Oracle Internet Directory, or `schemasync` for third-party LDAP directories.

4.9.3.4 Installing the LDAP Package

LDAP is included in a default Linux operating system installation.

If you did not perform a default Linux installation, and you intend to use Oracle scripts requiring LDAP, then use a package management system (`up2date`, `YaST`) for

your distribution to install a supported LDAP package for your distribution, and install any other required packages for that LDAP package.

4.9.4 Installing Oracle Messaging Gateway

Oracle Messaging Gateway is installed with Enterprise Edition of Oracle Database. However, you may require a CSD or Fix Packs.

If you require a CSD or Fix Packs for IBM WebSphere MQ, then see the following website for more information:

<http://www.ibm.com>

Note: Oracle Messaging Gateway does not support the integration of Advanced Queuing with TIBCO Rendezvous on IBM: Linux on System z.

Related Topics:

Oracle Database Advanced Queuing User's Guide

4.9.5 Installation Requirements for Programming Environments for Linux

Review the following section to install programming environments:

[Installation Requirements for Programming Environments for Linux x86-64](#)
(page 4-26)

Ensure that your system meets the requirements for the programming environment you want to configure:

4.9.5.1 Installation Requirements for Programming Environments for Linux x86-64

Ensure that your system meets the requirements for the programming environment you want to configure:

Table 4-9 Requirements for Programming Environments for Linux X86-64

Programming Environments	Support Requirements
Java Database Connectivity (JDBC) / Oracle Call Interface (OCI)	JDK 8 (Java SE Development Kit) with the JNDI extension with Oracle Java Database Connectivity.
Oracle C++ Oracle C++ Call Interface Pro*C/C++ Oracle XML Developer's Kit (XDK)	Intel C/C++ Compiler 12.0.5 or later, and the version of GNU C and C++ compilers listed in the software requirements section in this document for your platform. Oracle C++ Call Interface (OCCI) applications can be built only with Intel C++ Compiler 12.0.5 used with the standard template libraries of the gcc versions listed in the software requirements section in this document for your platform. Oracle XML Developer's Kit is supported with the same compilers as OCCI.
Pro*COBOL	<ul style="list-style-type: none"> Micro Focus Server Express 5.1 Micro Focus Visual COBOL for Eclipse 2.2 - Update 2

4.9.6 Installation Requirements for Web Browsers

Web browsers are required only if you intend to use Oracle Enterprise Manager Database Express and Oracle Enterprise Manager Cloud Control. Web browsers must support JavaScript, and the HTML 4.0 and CSS 1.0 standards.

For a list of browsers that meet these requirements see the Enterprise Manager certification matrix on My Oracle Support:

<https://support.oracle.com>

Related Topics:

Oracle Enterprise Manager Cloud Control Basic Installation Guide

4.10 Checking Kernel and Package Requirements for Linux

To check your kernel and packages to see if they meet minimum requirements for installation, perform the following steps:

1. To determine the distribution and version of Linux installed, enter one of the following commands:

```
# cat /etc/oracle-release
# cat /etc/redhat-release
# cat /etc/SuSE-release
# lsb_release -id
```

2. To determine if the required kernel errata is installed, enter the following command:

```
# uname -r
```

The following is an example of the output this command displays on an Oracle Linux 6 system:

```
2.6.39-100.7.1.el6uek.x86_64
```

Review the required errata level for your distribution. If the errata level is previous to the required minimum errata update, then obtain and install the latest kernel update from your Linux distributor.

3. To determine whether the required packages are installed, enter commands similar to the following:

```
# rpm -q package_name
```

Alternatively, if you require specific system architecture information, then enter the following command:

```
# rpm -qa --queryformat "%{NAME}-%{VERSION}-%{RELEASE} (%{ARCH})\n"
| grep package_name
```

You can also combine a query for multiple packages, and review the output for the correct versions. For example:

```
# rpm -q binutils compat-libstdc++ gcc glibc libaio libgcc libstdc++ \
make sysstat unixodbc
```

If a package is not installed, then install it from your Linux distribution media or download the required package version from your Linux distributor's website.

4.11 Installing the cvuqdisk RPM for Linux

If you do not use an Oracle Preinstallation RPM, and you want to use the Cluster Verification Utility, then you must install the `cvuqdisk` RPM.

Without `cvuqdisk`, Cluster Verification Utility cannot discover shared disks, and you receive the error message "Package `cvuqdisk` not installed" when you run Cluster Verification Utility. Use the `cvuqdisk rpm` for your hardware (for example, `x86_64`).

1. Locate the `cvuqdisk` RPM package, which is in the directory `rpm` on the Oracle Grid Infrastructure installation media. If you have already installed Oracle Grid Infrastructure, then it is located in the directory `grid_home/cv/rpm`.
2. Copy the `cvuqdisk` package to each node on the cluster. You should ensure that each node is running the same version of Linux.
3. Log in as `root`.
4. Use the following command to find if you have an existing version of the `cvuqdisk` package:

```
# rpm -qi cvuqdisk
```

If you have an existing version of `cvuqdisk`, then enter the following command to deinstall the existing version:

```
# rpm -e cvuqdisk
```

5. Set the environment variable `CVUQDISK_GRP` to point to the group that owns `cvuqdisk`, typically `oinstall`. For example:

```
# CVUQDISK_GRP=oinstall; export CVUQDISK_GRP
```

6. In the directory where you have saved the `cvuqdisk rpm`, use the command `rpm -iv package` to install the `cvuqdisk` package. For example:

```
# rpm -iv cvuqdisk-1.0.10-1.rpm
```

4.12 Reviewing HugePages Memory Allocation

Review this information if your operating system has HugePages enabled.

On Linux platform installations, Oracle recommends that you use HugePages to obtain the best performance for Oracle Databases. When you upgrade Oracle Grid Infrastructure and Oracle Databases on servers that have HugePages enabled, Oracle recommends that you review your HugePages memory allocation requirements.

GIMR and HugePages Memory

Oracle Grid Infrastructure installations include the Grid Infrastructure Management Repository (GIMR). When HugePages is configured on cluster member nodes, the GIMR system global area (SGA) is installed into HugePages memory. The GIMR SGA occupies up to 1 GB of HugePages memory. Oracle Grid Infrastructure starts up before Oracle Databases installed on the cluster.

If your cluster member node operating system memory allocations to HugePages are insufficient for the size of the SGAs for all of the Oracle Database instances on the cluster, then you may find that one or more of your Oracle Database SGAs are mapped to regular pages, instead of Huge Pages, which reduces expected performance. To avoid this issue, when you plan your upgrade, ensure that the

memory you reserve for HugePages is large enough to accommodate your memory requirements.

Allocate memory to HugePages large enough for all databases planned to run SGA on the cluster, and to accommodate the SGA for the Grid Infrastructure Management Repository.

4.13 Disabling Transparent HugePages

Oracle recommends that you disable Transparent HugePages before you start installation.

Transparent HugePages memory differs from standard HugePages memory because the kernel `khugepaged` thread allocates memory dynamically during runtime. Standard HugePages memory is pre-allocated at startup, and does not change during runtime.

Note: Although Transparent HugePages is disabled on UEK2 and later UEK kernels, Transparent HugePages may be enabled by default on your Linux system.

Transparent HugePages memory is enabled by default with Red Hat Enterprise Linux 6, Red Hat Enterprise Linux 7, SUSE 11, Oracle Linux 6, and Oracle Linux 7 with earlier releases of Oracle Linux with the Unbreakable Enterprise Kernel 2 (UEK2) kernels.

Transparent HugePages can cause memory allocation delays during runtime. To avoid performance issues, Oracle recommends that you disable Transparent HugePages on all Oracle Database servers. Oracle recommends that you instead use standard HugePages for enhanced performance.

To check if Transparent HugePages is enabled, run one of the following commands as the root user:

Red Hat Enterprise Linux kernels:

```
# cat /sys/kernel/mm/redhat_transparent_hugepage/enabled
```

Other kernels:

```
# cat /sys/kernel/mm/transparent_hugepage/enabled
```

The following is a sample output that shows Transparent HugePages are being used as the `[always]` flag is enabled.

```
[always] never
```

Note:

If Transparent HugePages is removed from the kernel, then neither `/sys/kernel/mm/transparent_hugepage` nor `/sys/kernel/mm/redhat_transparent_hugepage` files exist.

To disable Transparent HugePages:

1. For Oracle Linux 6 or earlier releases, add the following entry to the kernel boot line in the `/etc/grub.conf` file:

```
transparent_hugepage=never
```

For example:

```
title Oracle Linux Server (2.6.32-300.25.1.el6uek.x86_64)
  root (hd0,0)
  kernel /vmlinuz-2.6.32-300.25.1.el6uek.x86_64 ro root=LABEL=/
transparent_hugepage=never
  initrd /initramfs-2.6.32-300.25.1.el6uek.x86_64.img
```

The file name may vary for Oracle Linux 7 or later operating systems. Check your operating system documentation for the exact file name and the steps to disable Transparent HugePages.

For example, for Oracle Linux 7.3, the procedure to disable Transparent HugePages involves editing the `/etc/default/grub` file and then running the command `grub2-mkconfig`.

2. Restart the system to make the changes permanent.

4.14 Enabling the Name Service Cache Daemon

To allow Oracle Clusterware to better tolerate network failures with NAS devices or NFS mounts, enable the Name Service Cache Daemon (`nscd`).

To check to see if `nscd` is set to load when the system is restarted, enter the command `chkconfig --list nscd`. For example:

```
# chkconfig --list nscd
nscd                0:off 1:off 2:off 3:on  4:off 5:off 6:off
```

`nscd` is turned on for run level 3, and turned off for run level 5. The `nscd` should be turned on for both run level 3 and run level 5.

To change the configuration to ensure that `nscd` is on for both run level 3 and run level 5, enter the following command as `root`:

```
# chkconfig --level 35 nscd on
```

To start up `nscd` in the current session, enter the command as `root`:

```
# service nscd start
```

To restart `nscd` with the new setting, enter the following command as `root`:

```
# service nscd restart
```

For Oracle Linux 7, verify that `nscd` is enabled as follows:

```
systemctl --all |grep nscd
nscd.service loaded active running Name Service Cache Daemon
```

4.15 Verifying the Disk I/O Scheduler on Linux

For best performance for Oracle ASM, Oracle recommends that you use the Deadline I/O Scheduler.

Disk I/O schedulers reorder, delay, or merge requests for disk I/O to achieve better throughput and lower latency. Linux has multiple disk I/O schedulers available, including Deadline, Noop, Anticipatory, and Completely Fair Queuing (CFQ).

On each cluster node, enter the following command to verify that the Deadline disk I/O scheduler is configured for use:

```
# cat /sys/block/${ASM_DISK}/queue/scheduler
noop [deadline] cfq
```

In this example, the default disk I/O scheduler is Deadline and `ASM_DISK` is the Oracle Automatic Storage Management (Oracle ASM) disk device.

If the default disk I/O scheduler is not Deadline, then set it using a rules file:

1. Using a text editor, create a UDEV rules file for the Oracle ASM devices:

```
# vi /etc/udev/rules.d/60-oracle-schedulers.rules
```

2. Add the following line to the rules file and save it:

```
ACTION=="add|change", KERNEL=="sd[a-z]", ATTR{queue/rotational}=="0", ATTR{queue/scheduler}="deadline"
```

3. On clustered systems, copy the rules file to all other nodes on the cluster. For example:

```
$ scp 60-oracle-schedulers.rules root@node2:/etc/udev/rules.d/60-oracle-schedulers.rules
```

4. Load the rules file and restart the UDEV service. For example:

- a. Oracle Linux and Red Hat Enterprise Linux

```
# udevadm control --reload-rules
```

- b. SUSE Linux Enterprise Server

```
# /etc/init.d boot.udev restart
```

5. Verify that the disk I/O scheduler is set as Deadline.

4.16 Using Automatic SSH Configuration During Installation

To install Oracle software, configure secure shell (SSH) connectivity between all cluster member nodes.

Oracle Universal Installer (OUI) uses the `ssh` and `scp` commands during installation to run remote commands on and copy files to the other cluster nodes. You must configure SSH so that these commands do not prompt for a password.

Note:

Oracle configuration assistants use SSH for configuration operations from local to remote nodes. Oracle Enterprise Manager also uses SSH. RSH is no longer supported.

You can configure SSH from the OUI interface during installation for the user account running the installation. The automatic configuration creates passwordless SSH connectivity between all cluster member nodes. Oracle recommends that you use the automatic procedure if possible.

To enable the script to run, you must remove `stty` commands from the profiles of any existing Oracle software installation owners you want to use, and remove other

security measures that are triggered during a login, and that generate messages to the terminal. These messages, mail checks, and other displays prevent Oracle software installation owners from using the SSH configuration script that is built into OUI. If they are not disabled, then SSH must be configured manually before an installation can be run.

In rare cases, Oracle Clusterware installation may fail during the "AttachHome" operation when the remote node closes the SSH connection. To avoid this problem, set the following parameter in the SSH daemon configuration file `/etc/ssh/sshd_config` on all cluster nodes to set the timeout wait to unlimited:

```
LoginGraceTime 0
```

4.17 Setting Network Time Protocol for Cluster Time Synchronization

Use either the Oracle Cluster Time Synchronization or network time protocol (NTP) option for time synchronization.

Oracle Clusterware requires the same time zone environment variable setting on all cluster nodes. During installation, the installation process picks up the time zone (TZ) environment variable setting of the Grid installation owner on the node where Oracle Universal Installer (OUI) runs, and uses that time zone value on all nodes as the default TZ environment variable setting for all processes managed by Oracle Clusterware. The time zone default is used for databases, Oracle ASM, and any other managed processes. You have two options for time synchronization:

- An operating system configured network time protocol (NTP) such as `chronyd` or `ntpd`
- Oracle Cluster Time Synchronization Service

Oracle Cluster Time Synchronization Service is designed for organizations whose cluster servers are unable to access NTP services. If you use NTP, then the Oracle Cluster Time Synchronization daemon (`ctssd`) starts up in observer mode. If you do not have NTP daemons, then `ctssd` starts up in active mode and synchronizes time among cluster members without contacting an external time server.

Note:

- Before starting the installation of Oracle Grid Infrastructure, Oracle recommends that you ensure the clocks on all nodes are set to the same time.
 - By default, the NTP service available on Oracle Linux 7 and Red Hat Linux 7 is `chronyd` and the NTP service available on Oracle Linux 6 and Red Hat Linux 6 or lower is `ntpd`.
-
-

If you have NTP daemons on your server but you cannot configure them to synchronize time with a time server, and you want to use Cluster Time Synchronization Service to provide synchronization service in the cluster, then deactivate and deinstall the NTP.

Deactivating the `ntpd` Service

To deactivate the NTP service, you must stop the existing `ntpd` service, disable it from the initialization sequences, and remove the `ntp.conf` file.

Complete these step on Oracle Linux 6 and Red Hat Linux 6 or earlier, and Asianux systems:

1. Run the following commands as the `root` user:

```
# /sbin/service ntpd stop
# chkconfig ntpd off
# mv /etc/ntp.conf /etc/ntp.conf.org
```

2. Remove the following file:

```
/var/run/ntpd.pid
```

This file maintains the `pid` for the NTP daemon.

When the installer finds that the NTP protocol is not active, the Cluster Time Synchronization Service is installed in active mode and synchronizes the time across the nodes. If NTP is found configured, then the Cluster Time Synchronization Service is started in observer mode, and no active time synchronization is performed by Oracle Clusterware within the cluster.

Deactivating the chronyd Service

To deactivate the `chronyd` service, you must stop the existing `chronyd` service, and disable it from the initialization sequences.

Complete these step on Oracle Linux 7 and Red Hat Linux 7:

1. Run the following commands as the `root` user:

```
# systemctl stop chronyd
# systemctl disable chronyd
```

Confirming Oracle Cluster Time Synchronization Service After Installation

To confirm that `ctssd` is active after installation, enter the following command as the Grid installation owner:

```
$ crsctl check ctss
```

Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

[About Oracle Grid Infrastructure Network Configuration Options](#) (page 5-2)
Ensure that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

[Understanding Network Addresses](#) (page 5-3)
During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node.

[Network Interface Hardware Minimum Requirements](#) (page 5-6)
Review these requirements to ensure that you have the minimum network hardware technology for Oracle Grid Infrastructure clusters

[Private IP Interface Configuration Requirements](#) (page 5-7)
Requirements for private interfaces depend on whether you are using single or multiple Interfaces.

[IPv4 and IPv6 Protocol Requirements](#) (page 5-8)
Oracle Grid Infrastructure and Oracle RAC support the standard IPv6 address notations specified by RFC 2732 and global and site-local IPv6 addresses as defined by RFC 4193.

[Oracle Grid Infrastructure IP Name and Address Requirements](#) (page 5-9)
Review this information for Oracle Grid Infrastructure IP Name and Address requirements.

[Broadcast Requirements for Networks Used by Oracle Grid Infrastructure](#) (page 5-15)
Broadcast communications (ARP and UDP) must work properly across all the public and private interfaces configured for use by Oracle Grid Infrastructure.

[Multicast Requirements for Networks Used by Oracle Grid Infrastructure](#) (page 5-16)
For each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster.

[Domain Delegation to Grid Naming Service](#) (page 5-16)
If you are configuring Grid Naming Service (GNS) for a standard cluster, then before installing Oracle Grid Infrastructure you must configure

DNS to send to GNS any name resolution requests for the subdomain served by GNS.

[Configuration Requirements for Oracle Flex Clusters](#) (page 5-18)

Understand Oracle Flex Clusters and their configuration requirements.

[Grid Naming Service Cluster Configuration Example](#) (page 5-22)

Review this example to understand Grid Naming Service configuration.

[Manual IP Address Configuration Example](#) (page 5-24)

If you choose not to use GNS, then before installation you must configure public, virtual, and private IP addresses.

[Network Interface Configuration Options](#) (page 5-25)

During installation, you are asked to identify the planned use for each network adapter (or network interface) that Oracle Universal Installer (OUI) detects on your cluster node.

[Multiple Private Interconnects and Oracle Linux](#) (page 5-25)

Review this information for settings for the `rp_filter` parameter for Oracle RAC systems that use multiple NICs for the private interconnect.

5.1 About Oracle Grid Infrastructure Network Configuration Options

Ensure that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

Oracle Clusterware Networks

An Oracle Clusterware configuration requires at least two interfaces:

- A public network interface, on which users and application servers connect to access data on the database server
- A private network interface for internode communication.

You can configure a network interface to use either the IPv4 protocol, or the IPv6 protocol on a given network. If you use redundant network interfaces (bonded or teamed interfaces), then be aware that Oracle does not support configuring one interface to support IPv4 addresses and the other to support IPv6 addresses. You must configure network interfaces of a redundant interface pair with the same IP protocol.

All the nodes in the cluster must use the same IP protocol configuration. Either all the nodes use only IPv4, or all the nodes use only IPv6. You cannot have some nodes in the cluster configured to support only IPv6 addresses, and other nodes in the cluster configured to support only IPv4 addresses.

The VIP agent supports the generation of IPv6 addresses using the Stateless Address Autoconfiguration Protocol (RFC 2462), and advertises these addresses with GNS. Run the `srvctl config network` command to determine if Dynamic Host Configuration Protocol (DHCP) or stateless address autoconfiguration is being used.

Note: The Certify pages on My Oracle Support for the most up-to-date information about supported network protocols and hardware for Oracle RAC:

<https://support.oracle.com>

5.2 Understanding Network Addresses

During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node.

Identify each interface as a public or private interface, or as an interface that you do not want Oracle Grid Infrastructure or Oracle ASM to use. Public and virtual internet protocol (VIP) addresses are configured on public interfaces. Private addresses are configured on private interfaces.

Refer to the following sections for detailed information about each address type.

[About the Public IP Address](#) (page 5-3)

The public IP address is assigned dynamically using DHCP, or defined statically in a DNS or in a hosts file.

[About the Private IP Address](#) (page 5-3)

Oracle Clusterware uses interfaces marked as private for internode communication.

[About the Virtual IP Address](#) (page 5-4)

The virtual IP (VIP) address is registered in the grid naming service (GNS), or the DNS.

[About the Grid Naming Service \(GNS\) Virtual IP Address](#) (page 5-4)

The GNS virtual IP address is a static IP address configured in the DNS.

[About the SCAN](#) (page 5-5)

Oracle Database clients connect to the database using a single client access name (SCAN).

5.2.1 About the Public IP Address

The public IP address is assigned dynamically using DHCP, or defined statically in a DNS or in a hosts file.

The public IP address is assigned dynamically using Dynamic Host Configuration Protocol (DHCP), or defined statically in a Domain Name System (DNS) or in a hosts file. It uses the public interface (the interface with access available to clients). The public IP address is the primary address for a cluster member node, and should be the address that resolves to the name returned when you enter the command hostname.

If you configure IP addresses manually, then avoid changing host names after you complete the Oracle Grid Infrastructure installation, including adding or deleting domain qualifications. A node with a new host name is considered a new host, and must be added to the cluster. A node under the old name appears to be down until it is removed from the cluster.

5.2.2 About the Private IP Address

Oracle Clusterware uses interfaces marked as private for internode communication.

Each cluster node must have an interface that you identify during installation as a private interface. Private interfaces must have addresses configured for the interface itself, but no additional configuration is required. Oracle Clusterware uses the interfaces you identify as private for the cluster interconnect. If you identify multiple interfaces during information for the private network, then Oracle Clusterware configures them with Redundant Interconnect Usage. Any interface that you identify

as private must be on a subnet that connects to every node of the cluster. Oracle Clusterware uses all the interfaces you identify for use as private interfaces.

For the private interconnects, because of Cache Fusion and other traffic between nodes, Oracle strongly recommends using a physically separate, private network. If you configure addresses using a DNS, then you should ensure that the private IP addresses are reachable only by the cluster nodes.

You can choose multiple interconnects either during installation or postinstallation using the `oifcfg setif` command.

After installation, if you modify the interconnect for Oracle Real Application Clusters (Oracle RAC) with the `CLUSTER_INTERCONNECTS` initialization parameter, then you must change the interconnect to a private IP address, on a subnet that is not used with a public IP address, nor marked as a public subnet by `oifcfg`. Oracle does not support changing the interconnect to an interface using a subnet that you have designated as a public subnet.

You should not use a firewall on the network with the private network IP addresses, because this can block interconnect traffic.

5.2.3 About the Virtual IP Address

The virtual IP (VIP) address is registered in the grid naming service (GNS), or the DNS.

Select an address for your VIP that meets the following requirements:

- The IP address and host name are currently unused (it can be registered in a DNS, but should not be accessible by a ping command)
- The VIP is on the same subnet as your public interface

If you are not using Grid Naming Service (GNS), then determine a virtual host name for each node. A virtual host name is a public node name that reroutes client requests sent to the node if the node is down. Oracle Database uses VIPs for client-to-database connections, so the VIP address must be publicly accessible. Oracle recommends that you provide a name in the format `hostname-vip`. For example: `myclstr2-vip`.

5.2.4 About the Grid Naming Service (GNS) Virtual IP Address

The GNS virtual IP address is a static IP address configured in the DNS.

The DNS delegates queries to the GNS virtual IP address, and the GNS daemon responds to incoming name resolution requests at that address. Within the subdomain, the GNS uses multicast Domain Name Service (mDNS), included with Oracle Clusterware, to enable the cluster to map host names and IP addresses dynamically as nodes are added and removed from the cluster, without requiring additional host configuration in the DNS.

To enable GNS, you must have your network administrator provide a set of IP addresses for a subdomain assigned to the cluster (for example, `grid.example.com`), and delegate DNS requests for that subdomain to the GNS virtual IP address for the cluster, which GNS serves. DHCP provides the set of IP addresses to the cluster; DHCP must be available on the public network for the cluster.

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about GNS

5.2.5 About the SCAN

Oracle Database clients connect to the database using a single client access name (SCAN).

The SCAN and its associated IP addresses provide a stable name for clients to use for connections, independent of the nodes that make up the cluster. SCAN addresses, virtual IP addresses, and public IP addresses must all be on the same subnet.

The SCAN is a virtual IP name, similar to the names used for virtual IP addresses, such as `node1-vip`. However, unlike a virtual IP, the SCAN is associated with the entire cluster, rather than an individual node, and associated with multiple IP addresses, not just one address.

The SCAN resolves to multiple IP addresses reflecting multiple listeners in the cluster that handle public client connections. When a client submits a request, the SCAN listener listening on a SCAN IP address and the SCAN port is made available to a client. Because all services on the cluster are registered with the SCAN listener, the SCAN listener replies with the address of the local listener on the least-loaded node where the service is currently being offered. Finally, the client establishes connection to the service through the listener on the node where service is offered. All of these actions take place transparently to the client without any explicit configuration required in the client.

During installation, listeners are created. These SCAN listeners listen on the SCAN IP addresses. The SCAN listeners are started on nodes determined by Oracle Clusterware. Oracle Net Services routes application requests to the least-loaded instance providing the service. Because the SCAN addresses resolve to the cluster, rather than to a node address in the cluster, nodes can be added to or removed from the cluster without affecting the SCAN address configuration. The SCAN listener also supports HTTP protocol for communication with Oracle XML Database (XDB).

The SCAN should be configured so that it is resolvable either by using Grid Naming Service (GNS) within the cluster, or by using Domain Name Service (DNS) resolution. For high availability and scalability, Oracle recommends that you configure the SCAN name so that it resolves to three IP addresses. At a minimum, the SCAN must resolve to at least one address.

If you specify a GNS domain, then the SCAN name defaults to `clustername-scan.cluster_name.GNS_domain`. Otherwise, it defaults to `clustername-scan.current_domain`. For example, if you start Oracle Grid Infrastructure installation from the server `node1`, the cluster name is `mycluster`, and the GNS domain is `grid.example.com`, then the SCAN Name is `mycluster-scan.mycluster.grid.example.com`.

Clients configured to use IP addresses for Oracle Database releases prior to Oracle Database 11g release 2 can continue to use their existing connection addresses; using SCAN is not required. When you upgrade to Oracle Clusterware 12c release 1 (12.1), the SCAN becomes available, and you should use the SCAN for connections to Oracle Database 11g release 2 or later databases. When an earlier release of Oracle Database is upgraded, it registers with the SCAN listeners, and clients can start using the SCAN to connect to that database. The database registers with the SCAN listener through the remote listener parameter in the `init.ora` file. The `REMOTE_LISTENER` parameter must be set to `SCAN:PORT`. Do not set it to a `TNSNAMES` alias with a single address for the SCAN, for example, using `HOST= SCAN_name`.

The SCAN is optional for most deployments. However, clients using Oracle Database 11g release 2 and later policy-managed databases using server pools must access the

database using the SCAN. This is required because policy-managed databases can run on different servers at different times, so connecting to a particular node by using the virtual IP address for a policy-managed database is not possible.

Provide SCAN addresses for client access to the cluster. These addresses must be configured as round robin addresses on the domain name service (DNS), if DNS is used. Oracle recommends that you supply three SCAN addresses.

Identify public and private interfaces. Oracle Universal Installer configures public interfaces for use by public and virtual IP addresses, and configures private IP addresses on private interfaces. The private subnet that the private interfaces use must connect all the nodes you intend to have as cluster members. The SCAN must be in the same subnet as the public interface.

Related Topics:

Oracle Real Application Clusters Administration and Deployment Guide

5.3 Network Interface Hardware Minimum Requirements

Review these requirements to ensure that you have the minimum network hardware technology for Oracle Grid Infrastructure clusters

Public Network for Each Node

Public networks provide access to clients for database services. Public networks must meet these minimum requirements:

- **Adapters:** Each node must have at least one public network adapter or network interface cards (NIC).

Oracle supports the use of link aggregations, bonded, trunked or teamed networks for improved bandwidth and high availability.
- **Protocol:** Each public interface must support TCP/IP.

Private Network for Each Node

Private networks (also called interconnects) are networks that only cluster member nodes can access. They use switches for connections. Private networks must meet these minimum requirements:

- **Adapters:** Each node must have at least one private network adapter or network interface cards (NIC).

Oracle recommends that you configure interconnects using Redundant Interconnect Usage, in which multiple network adapters are configured with addresses in the link-local range to provide highly available IP (HAIP) addresses for the interconnect. You can configure Redundant Interconnect Usage either during installation, or after installation by using Oracle Interface Configuration Tool (OIFCFG), to provide improved bandwidth and high availability.

Oracle also supports the use of link aggregations, bonded, trunked or teamed networks for improved bandwidth and high availability.
- **Protocol:** User datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP, or Reliable Datagram Sockets (RDS) with Infiniband

Switches: You must use switches for interconnects that support TCP/IP. Oracle recommends that you use dedicated switches. The minimum switch speed is 1 Gigabit Ethernet.

Local Area Network Technology

Oracle does not support token-rings or crossover cables for the interconnect. Oracle supports Jumbo Frames and Infiniband. When you use Infiniband on the interconnect, Oracle supports using the RDS protocol.

If you have a shared Ethernet VLAN deployment, with shared physical adapter, ensure that you apply standard Ethernet design, deployment, and monitoring best practices to protect against cluster outages and performance degradation due to common shared Ethernet switch network events

Storage Networks

Oracle Automatic Storage Management and Oracle Real Application Clusters require network-attached storage.

Oracle Automatic Storage Management (Oracle ASM): The network interfaces used for Oracle Clusterware files are also used for Oracle ASM.

Third-party storage: Oracle recommends that you configure additional interfaces for storage.

5.4 Private IP Interface Configuration Requirements

Requirements for private interfaces depend on whether you are using single or multiple Interfaces.

Network Requirements for Single Interface Private Network Clusters

- Each node's private interface for interconnects must be on the same subnet.
- The subnet must connect to every node of the cluster.

For example, if the private interfaces have a subnet mask of 255.255.255.0, then your private network is in the range 192.168.0.0--192.168.0.255, and your private addresses must be in the range of 192.168.0.[0-255]. If the private interfaces have a subnet mask of 255.255.0.0, then your private addresses can be in the range of 192.168.[0-255].[0-255]

- Both IPv4 and IPv6 addresses are supported.

Network Requirements for Redundant Interconnect Usage Clusters

With Redundant Interconnect Usage, you can identify multiple interfaces to use for the cluster private network, without the need of using bonding or other technologies.

When you define multiple interfaces, Oracle Clusterware creates from one to four highly available IP (HAIP) addresses. Oracle RAC and Oracle Automatic Storage Management (Oracle ASM) instances use these interface addresses to ensure highly available, load-balanced interface communication between nodes. The installer enables Redundant Interconnect Usage to provide a high availability private network. By default, Oracle Grid Infrastructure software uses all of the HAIP addresses for private network communication, providing load-balancing across the set of interfaces you identify for the private network. If a private interconnect interface fails or become non-

communicative, then Oracle Clusterware transparently moves the corresponding HAIP address to one of the remaining functional interfaces.

- Each private interface should be on a different subnet.
- Each cluster member node must have an interface on each private interconnect subnet, and these subnets must connect to every node of the cluster.

For example, you can have private networks on subnets 192.168.0 and 10.0.0, but each cluster member node must have an interface connected to the 192.168.0 and 10.0.0 subnets.

- Endpoints of all designated interconnect interfaces must be completely reachable on the network. There should be no node that is not connected to every private network interface.

You can test if an interconnect interface is reachable using `ping`.

- You can use IPv4 and IPv6 addresses for the interfaces with Oracle Clusterware Redundant interconnects.

Note:

During installation, you can define up to four interfaces for the private network. The number of HAIP addresses created during installation is based on both physical and logical interfaces configured for the network adapter. After installation, you can define additional interfaces. If you define more than four interfaces as private network interfaces, then be aware that Oracle Clusterware activates only four of the interfaces at a time. However, if one of the four active interfaces fails, then Oracle Clusterware transitions the HAIP addresses configured to the failed interface to one of the reserve interfaces in the defined set of private interfaces.

See Also:

Oracle Clusterware Administration and Deployment Guide for more information about HAIP addresses

5.5 IPv4 and IPv6 Protocol Requirements

Oracle Grid Infrastructure and Oracle RAC support the standard IPv6 address notations specified by RFC 2732 and global and site-local IPv6 addresses as defined by RFC 4193.

Configuring Public VIPs

Cluster member node interfaces can be configured to use IPv4, IPv6, or both types of Internet protocol addresses. During installation, you can configure VIPs for a given public network as IPv4 or IPv6 types of addresses. You can configure an IPv6 cluster by selecting VIP and SCAN names that resolve to addresses in an IPv6 subnet for the cluster, and selecting that subnet as public during installation. After installation, you can also configure cluster member nodes with a mixture of IPv4 and IPv6 addresses.

If you install using static virtual IP (VIP) addresses in an IPv4 cluster, then the VIP names you supply during installation should resolve only to IPv4 addresses. If you

install using static IPv6 addresses, then the VIP names you supply during installation should resolve only to IPv6 addresses.

During installation, you cannot configure the cluster with VIP and SCAN names that resolve to both IPv4 and IPv6 addresses. You cannot configure VIPs and SCANS on some cluster member nodes to resolve to IPv4 addresses, and VIPs and SCANS on other cluster member nodes to resolve to IPv6 addresses. Oracle does not support this configuration.

Configuring Private IP Interfaces (Interconnects)

You can configure the private network either as an IPv4 network or IPv6 network.

Redundant Network Interfaces

If you configure redundant network interfaces for a public or VIP node name, then configure both interfaces of a redundant pair to the same address protocol. Also ensure that private IP interfaces use the same IP protocol. Oracle does not support names using redundant interface configurations with mixed IP protocols. You must configure both network interfaces of a redundant pair with the same IP protocol.

GNS or Multi-Cluster Addresses

Oracle Grid Infrastructure supports IPv4 DHCP addresses, and IPv6 addresses configured with the Stateless Address Autoconfiguration protocol, as described in RFC 2462.

Note: Link-local and site-local IPv6 addresses as defined in RFC 1884 are not supported.

5.6 Oracle Grid Infrastructure IP Name and Address Requirements

Review this information for Oracle Grid Infrastructure IP Name and Address requirements.

For small clusters, you can use a static configuration of IP addresses. For large clusters, manually maintaining the large number of required IP addresses becomes too cumbersome. Use Oracle Grid Naming Service with large clusters to ease network administration costs.

[About Oracle Grid Infrastructure Name Resolution Options](#) (page 5-10)

Before starting the installation, you must have at least two interfaces configured on each node: One for the private IP address and one for the public IP address.

[Cluster Name and SCAN Requirements](#) (page 5-11)

Review this information before selecting the cluster name and SCAN.

[IP Name and Address Requirements For Grid Naming Service \(GNS\)](#) (page 5-11)

Review this information for IP name and address requirements for Grid Naming Service (GNS).

[IP Name and Address Requirements For Multi-Cluster GNS](#) (page 5-11)

Multi-cluster GNS differs from standard GNS in that Multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

[IP Name and Address Requirements for Manual Configuration of Cluster](#)
(page 5-13)

For Oracle Flex Clusters and Application Clusters, configure static cluster node names and addresses if you do not enable GNS.

[Confirming the DNS Configuration for SCAN](#) (page 5-15)

Use the nslookup command to confirm that the DNS is correctly associating the SCAN with the addresses.

5.6.1 About Oracle Grid Infrastructure Name Resolution Options

Before starting the installation, you must have at least two interfaces configured on each node: One for the private IP address and one for the public IP address.

During installation, you are asked to identify the planned use for each network interface that Oracle Universal Installer (OUI) detects on your cluster node. Identify each interface as a public or private interface, or as an interface that you do not want Oracle Grid Infrastructure or Oracle ASM to use. Public and virtual internet protocol (VIP) addresses are configured on public interfaces. Private addresses are configured on private interfaces.

Configure IP addresses with one of the following options:

Dynamic IP address assignment using Multi-cluster or standard Oracle Grid Naming Service (GNS)

If you select this option, then network administrators delegate a subdomain to be resolved by GNS (standard or multicluster). Requirements for GNS are different depending on whether you choose to configure GNS with zone delegation (resolution of a domain delegated to GNS), or without zone delegation (a GNS virtual IP address without domain delegation).

For GNS with zone delegation:

- For IPv4, a DHCP service running on the public network the cluster uses
- For IPv6, an autoconfiguration service running on the public network the cluster uses
- Enough addresses on the DHCP server to provide one IP address for each node, and three IP addresses for the cluster used by the Single Client Access Name (SCAN) for the cluster

Use an existing GNS configuration

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), a single GNS instance can be used by multiple clusters. To use GNS for multiple clusters, the DNS administrator must have delegated a zone for use by GNS. Also, there must be an instance of GNS started somewhere on the network and the GNS instance must be accessible (not blocked by a firewall). All of the node names registered with the GNS instance must be unique.

Static IP address assignment using DNS or host file resolution

If you select this option, then network administrators assign a fixed IP address for each physical host name in the cluster and for IPs for the Oracle Clusterware managed VIPs. In addition, either domain name server (DNS) based static name resolution is used for each node, or host files for both the clusters and clients have to be updated, resulting in limited SCAN functionality. Selecting this option requires that you request network administration updates when you modify the cluster.

For GNS without zone delegation, configure a GNS virtual IP address (VIP) for the cluster. To enable Oracle Flex Cluster, you must at least configure a GNS virtual IP address.

5.6.2 Cluster Name and SCAN Requirements

Review this information before selecting the cluster name and SCAN.

Cluster Name and SCAN Requirements

Cluster Name must meet the following requirements:

- The cluster name is case-insensitive, must be unique across your enterprise, must be at least one character long and no more than 15 characters in length, must be alphanumeric, cannot begin with a numeral, and may contain hyphens (-). Underscore characters (_) are not allowed.
- The SCAN and cluster name are entered in separate fields during installation, so cluster name requirements do not apply to the name used for the SCAN, and the SCAN can be longer than 15 characters. If you enter a domain with the SCAN name, and you want to use GNS with zone delegation, then the domain must be the GNS domain.

Note: Select your cluster name carefully. After installation, you can only change the cluster name by reinstalling Oracle Grid Infrastructure.

5.6.3 IP Name and Address Requirements For Grid Naming Service (GNS)

Review this information for IP name and address requirements for Grid Naming Service (GNS).

IP Name and Address Requirements For Grid Naming Service (GNS)

If you enable Grid Naming Service (GNS), then name resolution requests to the cluster are delegated to the GNS, which is listening on the GNS virtual IP address. The domain name server (DNS) must be configured to delegate resolution requests for cluster names (any names in the subdomain delegated to the cluster) to the GNS. When a request comes to the domain, GNS processes the requests and responds with the appropriate addresses for the name requested. To use GNS, you must specify a static IP address for the GNS VIP address.

Note: The following restrictions apply to vendor configurations on your system:

- You cannot use GNS with another multicast DNS. To use GNS, disable any third party mDNS daemons on your system.
-
-

5.6.4 IP Name and Address Requirements For Multi-Cluster GNS

Multi-cluster GNS differs from standard GNS in that Multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

[About Multi-Cluster GNS Networks](#) (page 5-12)

The general requirements for multi-cluster GNS are similar to those for standard GNS. Multi-cluster GNS differs from standard GNS in that multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

[Configuring GNS Server Clusters](#) (page 5-12)

Review these requirements to configure GNS server clusters.

[Configuring GNS Client Clusters](#) (page 5-13)

To configure a GNS client cluster, check to ensure all of the following requirements are completed.

[Creating and Using a GNS Client Data File](#) (page 5-13)

Generate a GNS client data file and copy the file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

5.6.4.1 About Multi-Cluster GNS Networks

The general requirements for multi-cluster GNS are similar to those for standard GNS. Multi-cluster GNS differs from standard GNS in that multi-cluster GNS provides a single networking service across a set of clusters, rather than a networking service for a single cluster.

Requirements for Multi-Cluster GNS Networks

To provide networking service, multi-cluster Grid Naming Service (GNS) is configured using DHCP addresses, and name advertisement and resolution is carried out with the following components:

- The GNS server cluster performs address resolution for GNS client clusters. A GNS server cluster is the cluster where multi-cluster GNS runs, and where name resolution takes place for the subdomain delegated to the set of clusters.
- GNS client clusters receive address resolution from the GNS server cluster. A GNS client cluster is a cluster that advertises its cluster member node names using the GNS server cluster.
- If you choose to use GNS, then the GNS configured at the time of installation is the primary. A secondary GNS for high availability can be configured at a later time.

5.6.4.2 Configuring GNS Server Clusters

Review these requirements to configure GNS server clusters.

To configure a GNS server cluster, check to ensure all of the following requirements are completed:

- Your network administrators must have delegated a subdomain to GNS for resolution.
- Before installation, create a static IP address for the GNS VIP address, and provide a subdomain that your DNS servers delegate to that static GNS IP address for resolution.

5.6.4.3 Configuring GNS Client Clusters

To configure a GNS client cluster, check to ensure all of the following requirements are completed.

- A GNS server instance must be running on your network, and it must be accessible (for example, not blocked by a firewall).
- All of the node names in the GNS domain must be unique; address ranges and cluster names must be unique for both GNS server and GNS client clusters.
- You must have a GNS client data file that you generated on the GNS server cluster, so that the GNS client cluster has the information needed to delegate its name resolution to the GNS server cluster, and you must have copied that file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

5.6.4.4 Creating and Using a GNS Client Data File

Generate a GNS client data file and copy the file to the GNS client cluster member node on which you are running the Oracle Grid Infrastructure installation.

On a GNS server cluster member, run the following command, where *path_to_file* is the name and path location of the GNS client data file you create:

```
srvctl export gns -clientdata path_to_file -role client
```

For example:

```
$ srvctl export gns -clientdata /home/grid/gns_client_data -role client
```

Copy the GNS Client data file to a secure path on the GNS Client node where you run the GNS Client cluster installation. The Oracle installation user must have permissions to access that file. Oracle recommends that no other user is granted permissions to access the GNS Client data file. During installation, you are prompted to provide a path to that file.

```
srvctl add gns -clientdata path_to_file
```

For example:

```
$ srvctl add gns -clientdata /home/grid/gns_client_data
```

See [Oracle Clusterware Administration and Deployment Guide](#) for more information about GNS server and GNS client administration

5.6.5 IP Name and Address Requirements for Manual Configuration of Cluster

For Oracle Flex Clusters and Application Clusters, configure static cluster node names and addresses if you do not enable GNS.

IP Address Requirements for Static Clusters

Public and virtual IP names must conform with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_").

Oracle Clusterware manages private IP addresses in the private subnet on interfaces you identify as private during the installation interview.

Public IP Address Requirements

The cluster must have a public IP address for each node, with the following characteristics:

- Static IP address
- Configured before installation for each node, and resolvable to that node before installation
- On the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster

Virtual IP Address Requirements

The cluster must have a virtual IP address for each node, with the following characteristics:

- Static IP address
- Configured before installation for each node, but not currently in use
- On the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster

Single Client Access Name Requirements

The cluster must have a Single Client Access Name (SCAN) for the cluster, with the following characteristics:

- Three static IP addresses configured on the domain name server (DNS) before installation so that the three IP addresses are associated with the name provided as the SCAN, and all three addresses are returned in random order by the DNS to the requestor
- Configured before installation in the DNS to resolve to addresses that are not currently in use
- Given addresses on the same subnet as all other public IP addresses, VIP addresses, and SCAN addresses in the cluster
- Given a name that does not begin with a numeral, and that conforms with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_")

Private IP Address Requirements

The cluster must have a private IP address for each node, with the following characteristics:

- Static IP address
- Configured before installation, but on a separate, private network, with its own subnet, that is not resolvable except by other cluster member nodes

The SCAN is a name used to provide service access for clients to the cluster. Because the SCAN is associated with the cluster as a whole, rather than to a particular node, the SCAN makes it possible to add or remove nodes from the cluster without needing to reconfigure clients. It also adds location independence for the databases, so that

client configuration does not have to depend on which nodes are running a particular database. Clients can continue to access the cluster in the same way as with previous releases, but Oracle recommends that clients accessing the cluster use the SCAN.

Note:

The SCAN and cluster name are entered in separate fields during installation, so cluster name requirements do not apply to the SCAN name.

Oracle strongly recommends that you do not configure SCAN VIP addresses in the hosts file. Use DNS resolution for SCAN VIPs. If you use the hosts file to resolve SCANS, then the SCAN can resolve to one IP address only.

Configuring SCANS in a DNS or a hosts file is the only supported configuration. Configuring SCANS in a Network Information Service (NIS) is not supported.

5.6.6 Confirming the DNS Configuration for SCAN

Use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses.

The following example shows how to use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses:

```
root@node1]$ nslookup mycluster-scan
Server:          dns.example.com
Address:         192.0.2.001

Name:   mycluster-scan.example.com
Address: 192.0.2.201
Name:   mycluster-scan.example.com
Address: 192.0.2.202
Name:   mycluster-scan.example.com
Address: 192.0.2.203
```

After installation, when a client sends a request to the cluster, the Oracle Clusterware SCAN listeners redirect client requests to servers in the cluster.

Oracle strongly recommends that you do not configure SCAN VIP addresses in the hosts file. Use DNS resolution for SCAN VIPs. If you use the hosts file to resolve SCANS, then the SCAN can resolve to one IP address only.

Configuring SCANS in a DNS or a hosts file is the only supported configuration. Configuring SCANS in a Network Information Service (NIS) is not supported.

5.7 Broadcast Requirements for Networks Used by Oracle Grid Infrastructure

Broadcast communications (ARP and UDP) must work properly across all the public and private interfaces configured for use by Oracle Grid Infrastructure.

The broadcast must work across any configured VLANs as used by the public or private interfaces.

When configuring public and private network interfaces for Oracle RAC, you must enable Address Resolution Protocol (ARP). Highly Available IP (HAIP) addresses do

not require ARP on the public network, but for VIP failover, you need to enable ARP. Do not configure NOARP.

5.8 Multicast Requirements for Networks Used by Oracle Grid Infrastructure

For each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster.

Multicast Requirements for Networks Used by Oracle Grid Infrastructure

Multicasting is required on the private interconnect. For this reason, at a minimum, you must enable multicasting for the cluster:

- Across the broadcast domain as defined for the private interconnect
- On the IP address subnet ranges 224.0.0.0/24 and optionally 230.0.1.0/24

You do not need to enable multicast communications across routers.

5.9 Domain Delegation to Grid Naming Service

If you are configuring Grid Naming Service (GNS) for a standard cluster, then before installing Oracle Grid Infrastructure you must configure DNS to send to GNS any name resolution requests for the subdomain served by GNS.

The subdomain that GNS serves represents the cluster member nodes.

[Choosing a Subdomain Name for Use with Grid Naming Service](#) (page 5-16)

To implement GNS, your network administrator must configure the DNS to set up a domain for the cluster, and delegate resolution of that domain to the GNS VIP.

[Configuring DNS for Cluster Domain Delegation to Grid Naming Service](#) (page 5-17)

If you plan to use Grid Naming Service (GNS) with a delegated domain, then before Oracle Grid Infrastructure installation, configure your domain name server (DNS) to send to GNS name resolution requests for the subdomain GNS serves, which are the cluster member nodes.

5.9.1 Choosing a Subdomain Name for Use with Grid Naming Service

To implement GNS, your network administrator must configure the DNS to set up a domain for the cluster, and delegate resolution of that domain to the GNS VIP.

Requirements for Choosing a Subdomain Name for Use with GNS

You can use a separate domain, or you can create a subdomain of an existing domain for the cluster. The subdomain name can be any supported DNS name such as `sales-cluster.rac.com`.

Oracle recommends that the subdomain name is distinct from your corporate domain. For example, if your corporate domain is `mycorp.example.com`, the subdomain for GNS might be `rac-gns.mycorp.example.com`.

If the subdomain is not distinct, then it should be for the exclusive use of GNS. For example, if you delegate the subdomain `mydomain.example.com` to GNS, then there should be no other domains that share it such as `lab1.mydomain.example.com`.

5.9.2 Configuring DNS for Cluster Domain Delegation to Grid Naming Service

If you plan to use Grid Naming Service (GNS) with a delegated domain, then before Oracle Grid Infrastructure installation, configure your domain name server (DNS) to send to GNS name resolution requests for the subdomain GNS serves, which are the cluster member nodes.

GNS domain delegation is mandatory with dynamic public networks (DHCP, autoconfiguration). GNS domain delegation is not required with static public networks (static addresses, manual configuration).

The following is an overview of the steps to be performed for domain delegation. Your actual procedure may be different from this example.

Configure the DNS to send GNS name resolution requests using delegation:

1. In the DNS, create an entry for the GNS virtual IP address, where the address uses the form `gns-server.clustername.domainname`. For example, where the cluster name is `mycluster`, and the domain name is `example.com`, and the IP address is `192.0.2.1`, create an entry similar to the following:

```
mycluster-gns-vip.example.com A 192.0.2.1
```

The address you provide must be routable.

2. Set up forwarding of the GNS subdomain to the GNS virtual IP address, so that GNS resolves addresses to the GNS subdomain. To do this, create a BIND configuration entry similar to the following for the delegated domain, where `cluster01.example.com` is the subdomain you want to delegate:

```
cluster01.example.com NS mycluster-gns-vip.example.com
```

3. When using GNS, you must configure `resolve.conf` on the nodes in the cluster (or the file on your system that provides resolution information) to contain name server entries that are resolvable to corporate DNS servers. The total timeout period configured—a combination of options attempts (retries) and options timeout (exponential backoff)—should be less than 30 seconds. For example, where `xxx.xxx.xxx.42` and `xxx.xxx.xxx.15` are valid name server addresses in your network, provide an entry similar to the following in `/etc/resolv.conf`:

```
options attempts: 2
options timeout: 1
```

```
search cluster01.example.com example.com
nameserver xxx.xxx.xxx.42
nameserver xxx.xxx.xxx.15
```

`/etc/nsswitch.conf` controls name service lookup order. In some system configurations, the Network Information System (NIS) can cause problems with SCAN address resolution. Oracle recommends that you place the NIS entry at the end of the search list. For example:

```
/etc/nsswitch.conf
hosts:    files  dns  nis
```

Be aware that use of NIS is a frequent source of problems when doing cable pull tests, as host name and user name resolution can fail.

5.10 Configuration Requirements for Oracle Flex Clusters

Understand Oracle Flex Clusters and their configuration requirements.

[Understanding Oracle Flex Clusters](#) (page 5-18)

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure cluster configurations are Oracle Flex Clusters deployments.

[About Oracle Flex ASM Clusters Networks](#) (page 5-19)

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), as part of an Oracle Flex Cluster installation, Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services.

[General Requirements for Oracle Flex Cluster Configuration](#) (page 5-21)

Review this information about network requirements for Oracle Flex Cluster configuration.

[Oracle Flex Cluster DHCP-Assigned Virtual IP \(VIP\) Addresses](#) (page 5-21)

Configure cluster node VIP names for both Hub and Leaf Nodes.

[Oracle Flex Cluster Manually-Assigned Addresses](#) (page 5-21)

Review this information to manually assign cluster node VIP names for both Hub and Leaf Nodes.

5.10.1 Understanding Oracle Flex Clusters

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure cluster configurations are Oracle Flex Clusters deployments.

Oracle Grid Infrastructure installed in an Oracle Flex Cluster configuration is a scalable, dynamic, robust network of nodes. Oracle Flex Clusters provide a platform for Oracle Real Application Clusters databases with large numbers of nodes, to support massive parallel query operations. Oracle Flex Clusters also provide a platform for other service deployments that require coordination and automation for high availability.

All nodes in an Oracle Flex Cluster belong to a single Oracle Grid Infrastructure cluster. This architecture centralizes policy decisions for deployment of resources based on application needs, to account for various service levels, loads, failure responses, and recovery.

Oracle Flex Clusters contain two types of nodes arranged in a hub and spoke architecture: Hub Nodes and Leaf Nodes. The number of Hub Nodes in an Oracle Flex Cluster can be as many as 64. The number of Leaf Nodes can be many more. Hub Nodes and Leaf Nodes can host different types of applications and perform parallel query operations.

Hub Nodes in Oracle Flex Clusters are tightly connected, and have direct access to shared storage. In an Oracle Flex Cluster configuration, Hub Nodes can also provide storage service for one or more Leaf Nodes. Three Hub Nodes act as I/O Servers for storage access requests from Leaf Nodes. If a Hub Node acting as an I/O Server becomes unavailable, then Oracle Grid Infrastructure starts another I/O Server on another Hub Node.

Leaf Nodes in Oracle Flex Clusters do not require direct access to shared storage, but instead request data through Hub Nodes. Hub Nodes can run in an Oracle Flex Cluster configuration without having any Leaf Nodes as cluster member nodes, but Leaf Nodes must be members of a cluster with a pool of Hub Nodes.

Oracle RAC database instances running on Leaf Nodes are referred to as far Oracle ASM client instances. Oracle ASM metadata is never sent to the far client database instance. Instead, the far Oracle ASM client database sends the I/O requests to I/O Server instances running on Hub Nodes over the Oracle Flex ASM network.

You configure servers for Hub Node and Leaf Node roles. You can designate servers for manual or automatic configuration.

If you select manual configuration, then you must designate each node in your cluster as a Hub Node or a Leaf Node. Each role requires different access to storage. To be eligible for the Hub Node role, a server must have direct access to storage. To be eligible for the Leaf Node role, a server may have access to direct storage, but it does not require direct access, because leaf nodes access storage as clients through Hub Nodes.

If you select automatic configuration of roles, then cluster nodes that have access to storage and join are configured as Hub Nodes, up to the number that you designate as your target. Other nodes that do not have access to storage or that join the cluster after that target number is reached join the cluster as Leaf Nodes. Nodes are configured as needed to provide Hub Nodes configured with Local or Near ASM to provide storage client services, and Leaf Nodes that are configured with direct access to Oracle ASM disks can be reconfigured as needed to become Hub Nodes. Oracle recommends that you select automatic configuration of Hub and Leaf node roles.

About Reader Nodes

You can use Leaf Nodes to host Oracle RAC database instances that run in read-only mode, which become reader nodes. You can optimize these nodes for parallel query operations by provisioning nodes with a large amount of memory so that data is cached in the Leaf Node.

A Leaf Node sends periodic heartbeat messages to its associated Hub Node, which is different from the heartbeat messages that occur between Hub Nodes. During planned shutdown of the Hub Nodes, a Leaf Node attempts to connect to another Hub Node, unless the Leaf Node is connected to only one Hub Node. If the Hub Node is evicted, then the Leaf Node is also evicted from the cluster.

5.10.2 About Oracle Flex ASM Clusters Networks

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), as part of an Oracle Flex Cluster installation, Oracle ASM is configured within Oracle Grid Infrastructure to provide storage services.

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. Many Oracle ASM instances can be clustered to support numerous database clients. Each Oracle Flex ASM cluster has its own name that is globally unique within the enterprise.

You can consolidate all the storage requirements into a single set of disk groups. All these disk groups are managed by a small set of Oracle ASM instances running in a single Oracle Flex Cluster.

Every Oracle Flex ASM cluster has one or more Hub Nodes on which Oracle ASM instances are running.

Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. Each network can be classified PUBLIC, ASM & PRIVATE, PRIVATE, or ASM.

The Oracle ASM network can be configured during installation, or configured or modified after installation.

About Oracle Flex ASM Cluster Configuration on Hub Nodes

Oracle Flex ASM cluster Hub Nodes can be configured with the following characteristics:

- Are similar to prior release Oracle Grid Infrastructure cluster member nodes, as all servers configured with the Hub Node role are peers.
- Have direct connections to the Oracle ASM disks.
- Run a Direct ASM client process.
- Run an Oracle ASM Filter Driver, part of whose function is to provide cluster fencing security for the Oracle Flex ASM cluster.
- Access the Oracle ASM disks as Hub Nodes only, where they are designated a Hub Node for that storage.
- Respond to service requests delegated to them through the global Oracle ASM listener configured for the Oracle Flex ASM cluster, which designates three of the Oracle Flex ASM cluster member Hub Node listeners as remote listeners for the Oracle Flex ASM cluster.
- Can provide database clients that are running on Hub nodes of the Oracle ASM cluster remote access to Oracle ASM for metadata, and allow database clients to perform block I/O operations directly to Oracle ASM disks. The hosts running the Oracle ASM server and the remote database client must both be Hub Nodes.

About Oracle Flex ASM Cluster Configuration on Leaf Nodes

Oracle Flex ASM cluster Leaf Nodes can be configured with the following characteristics:

- Use Indirect access to the Oracle ASM disks, where I/O is handled as a service for the client on a Hub Node.
- Submit disk service requests through the Oracle ASM network

About Oracle Flex ASM Cluster with Oracle IO Server (IOS) Configuration

An Oracle IO Server instance provides Oracle ASM file access for Oracle Database instances on nodes of Oracle Member Clusters that do not have connectivity to Oracle ASM managed disks. IOS enables you to configure Oracle Member Clusters on such nodes. On the storage cluster, the IO Server instance on each node opens up network ports to which clients send their I/O. The IO Server instance receives data packets from the client and performs the appropriate I/O to Oracle ASM disks similar to any other database client. On the client side, databases can use direct NFS (dNFS) to communicate with an IO Server instance. However, no client side configuration is required to use IO Server, so you are not required to provide a server IP address or any additional configuration information. On nodes and clusters that are configured to access Oracle ASM files through IO Server, the discovery of the Oracle IOS instance occurs automatically.

To install an Oracle Member Cluster, the administrator of the Oracle Domain Services Cluster creates an Oracle Member Cluster using a `crsctl` command that creates a Member Cluster Manifest file. During Oracle Grid Infrastructure installation, if you choose to install an Oracle Member Cluster, then the installer prompts you for the Member Cluster Manifest file. An attribute in the Member Cluster Manifest file specifies if the Oracle Member Cluster is expected to access Oracle ASM files through an IO Server instance.

See Also: *Oracle Automatic Storage Management Administrator's Guide*

5.10.3 General Requirements for Oracle Flex Cluster Configuration

Review this information about network requirements for Oracle Flex Cluster configuration.

Network Requirements for Oracle Flex Cluster Configuration

- You must use Grid Naming Service (GNS) with an Oracle Flex Cluster deployment.
- You must configure the GNS VIP as a static IP address for Hub Nodes.
- On Multi-cluster configurations, you must identify the GNS client data file location for Leaf Nodes. The GNS client data files are copied over from the GNS server before you start configuring a GNS client cluster.
- All public network addresses for both Hub Nodes and Leaf Nodes, whether assigned manually or automatically, must be in the same subnet range.
- All Oracle Flex Cluster addresses must be either static IP addresses, DHCP addresses assigned through DHCP (IPv4) or autoconfiguration addresses assigned through an autoconfiguration service (IPv6), registered in the cluster through GNS.
- When using GNS, you can also configure Leaf Nodes on both public and private networks, during installation. Leaf Nodes on public networks cannot use Oracle Clusterware services such as the public network resources and VIPs, or run listeners. After installation, you can configure network resources and listeners for the Leaf Nodes using SRVCTL commands.

5.10.4 Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses

Configure cluster node VIP names for both Hub and Leaf Nodes.

Requirements for DHCP-Assigned VIP Addresses

If you want to configure DHCP-assigned VIPs, then during installation, configure cluster node VIP names for both Hub and Leaf Nodes as follows:

- **Automatically Assigned Names:** Select the **Configure nodes Virtual IPs assigned by the Dynamic Networks** option to allow the installer to assign names to VIP addresses generated through DHCP automatically. Addresses are assigned through DHCP, and resolved by GNS. Oracle Clusterware sends DHCP requests with `client ID nodename-vip` and without a MAC address. You can verify the availability of DHCP addresses using the `cluvfy comp dhcp` command.

5.10.5 Oracle Flex Cluster Manually-Assigned Addresses

Review this information to manually assign cluster node VIP names for both Hub and Leaf Nodes.

Requirements for Manually-Assigned Addresses

If you choose to configure manually-assigned VIPs, then during installation, you must configure cluster node VIP names for both Hub and Leaf Nodes using one of the following options:

- **Manual Names:** Enter the host name and virtual IP name for each node manually, and select whether it is a Hub Node or a Leaf Node. The names you provide must resolve to addresses configured on the DNS. Names must conform with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_").
- **Automatically Assigned Names:** Enter string variables for values corresponding to host names that you have configured on the DNS. String variables allow you to assign a large number of names rapidly during installation. Configure addresses on the DNS with the following characteristics:
 - **Hostname prefix:** a prefix string used in each address configured on the DNS for use by cluster member nodes. For example: `mycloud`.
 - **Range:** A range of numbers to be assigned to the cluster member nodes, consisting of a starting node number and an ending node number, designating the end of the range. For example: `001` and `999`.
 - **Node name suffix:** A suffix added after the end of a range number to a public node name. For example: `nd`.
 - **VIP name suffix:** A suffix added after the end of a virtual IP node name. For example: `-vip`.

Syntax

You can create manual addresses using alphanumeric strings.

Example 5-1 Examples of Manually-Assigned Addresses

```
mycloud001nd; mycloud046nd; mycloud046-vip; mycloud348nd;  
mycloud784-vip
```

5.11 Grid Naming Service Cluster Configuration Example

Review this example to understand Grid Naming Service configuration.

To use GNS, you must specify a static IP address for the GNS VIP address, and you must have a subdomain configured on your DNS to delegate resolution for that subdomain to the static GNS IP address.

As nodes are added to the cluster, your organization's DHCP server can provide addresses for these nodes dynamically. These addresses are then registered automatically in GNS, and GNS provides resolution within the subdomain to cluster node addresses registered with GNS.

Because allocation and configuration of addresses is performed automatically with GNS, no further configuration is required. Oracle Clusterware provides dynamic network configuration as nodes are added to or removed from the cluster. The following example is provided only for information.

With IPv6 networks, the IPv6 auto configuration feature assigns IP addresses and no DHCP server is required.

With a two node cluster where you have defined the GNS VIP, after installation you might have a configuration similar to the following for a two-node cluster, where the cluster name is mycluster, the GNS parent domain is gns.example.com, the subdomain is cluster01.example.com, the 192.0.2 portion of the IP addresses represents the cluster public IP address subdomain, and 192.168 represents the private IP address subdomain:

Table 5-1 Grid Naming Service Cluster Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
GNS VIP	None	Selected by Oracle Clusterware	mycluster-gns-vip.example.com	virtual	192.0.2.1	Fixed by net administrator	DNS
Node 1 Public	Node 1	node1	node1	public	192.0.2.101	Fixed	GNS
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	virtual	192.0.2.104	DHCP	GNS
Node 1 Private	Node 1	node1	node1-priv	private	192.168.0.1	Fixed or DHCP	GNS
Node 2 Public	Node 2	node2	node2	public	192.0.2.102	Fixed	GNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	virtual	192.0.2.105	DHCP	GNS
Node 2 Private	Node 2	node2	node2-priv	private	192.168.0.2	Fixed or DHCP	GNS
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.201	DHCP	GNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.202	DHCP	GNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan.mycluster.cluster01.example.com	virtual	192.0.2.203	DHCP	GNS

5.12 Manual IP Address Configuration Example

If you choose not to use GNS, then before installation you must configure public, virtual, and private IP addresses.

Check that the default gateway can be accessed by a ping command. To find the default gateway, use the `route` command, as described in your operating system's help utility.

For example, with a two-node cluster where each node has one public and one private interface, and you have defined a SCAN domain address to resolve on your DNS to one of three IP addresses, you might have the configuration shown in the following table for your network interfaces:

Table 5-2 Manual Network Configuration Example

Identity	Home Node	Host Node	Given Name	Type	Address	Address Assigned By	Resolved By
Node 1 Public	Node 1	node1	node1	public	192.0.2.101	Fixed	DNS
Node 1 VIP	Node 1	Selected by Oracle Clusterware	node1-vip	virtual	192.0.2.104	Fixed	DNS and hosts file
Node 1 Private	Node 1	node1	node1-priv	private	192.168.0.1	Fixed	DNS and hosts file, or none
Node 2 Public	Node 2	node2	node2	public	192.0.2.102	Fixed	DNS
Node 2 VIP	Node 2	Selected by Oracle Clusterware	node2-vip	virtual	192.0.2.105	Fixed	DNS and hosts file
Node 2 Private	Node 2	node2	node2-priv	private	192.168.0.2	Fixed	DNS and hosts file, or none
SCAN VIP 1	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.201	Fixed	DNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.202	Fixed	DNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan	virtual	192.0.2.203	Fixed	DNS

You do not need to provide a private name for the interconnect. If you want name resolution for the interconnect, then you can configure private IP names in the hosts file or the DNS. However, Oracle Clusterware assigns interconnect addresses on the interface defined during installation as the private interface (eth1, for example), and to the subnet used for the private subnet.

The addresses to which the SCAN resolves are assigned by Oracle Clusterware, so they are not fixed to a particular node. To enable VIP failover, the configuration shown in the preceding table defines the SCAN addresses and the public and VIP addresses of both nodes on the same subnet, 192.0.2.

Note: All host names must conform to the RFC-952 standard, which permits alphanumeric characters. Host names using underscores ("_") are not allowed.

5.13 Network Interface Configuration Options

During installation, you are asked to identify the planned use for each network adapter (or network interface) that Oracle Universal Installer (OUI) detects on your cluster node.

Each NIC can be configured to perform only one of the following roles:

- Public
- Private
- Do Not Use

Network Interface Configuration Options

You must use the same private adapters for both Oracle Clusterware and Oracle RAC. The precise configuration you choose for your network depends on the size and use of the cluster you want to configure, and the level of availability you require. Network interfaces must be at least 1 GbE, with 10 GbE recommended. Alternatively, use InfiniBand for the interconnect.

If certified Network-attached Storage (NAS) is used for Oracle RAC and this storage is connected through Ethernet-based networks, then you must have a third network interface for NAS I/O. Failing to provide three separate interfaces in this case can cause performance and stability problems under load.

Redundant interconnect usage cannot protect network adapters used for public communication. If you require high availability or load balancing for public adapters, then use a third party solution. Typically, bonding, trunking or similar technologies can be used for this purpose.

You can enable redundant interconnect usage for the private network by selecting multiple network adapters to use as private adapters. Redundant interconnect usage creates a redundant interconnect when you identify more than one network adapter as private.

5.14 Multiple Private Interconnects and Oracle Linux

Review this information for settings for the `rp_filter` parameter for Oracle RAC systems that use multiple NICs for the private interconnect.

With Oracle Linux kernel 2.6.31, which also includes Oracle Unbreakable Enterprise Kernel 2.6.32, a bug has been fixed in the Reverse Path Filtering. As a consequence of this correction, Oracle RAC systems that use multiple NICs for the private

interconnect now require specific settings for the `rp_filter` parameter. This requirement also applies to all Exadata systems that are running Linux kernel 2.6.32 and above. Without these `rp_filter` parameter settings systems, interconnect packets can be blocked or discarded.

The `rp_filter` values set the Reverse Path filter to no filtering (0), to strict filtering (1), or to loose filtering (2). Set the `rp_filter` value for the private interconnects to either 0 or 2. Setting the private interconnect NIC to 1 can cause connection issues on the private interconnect. It is not considered unsafe to disable or relax this filtering, because the private interconnect should be on a private and isolated network.

For example, where `eth1` and `eth2` are the private interconnect NICs, and `eth0` is the public network NIC, set the `rp_filter` of the private address to 2 (loose filtering), the public address to 1 (strict filtering), using the following entries in `/etc/sysctl.conf`:

```
net.ipv4.conf.eth2.rp_filter = 2
net.ipv4.conf.eth1.rp_filter = 2
net.ipv4.conf.eth0.rp_filter = 1
```

My Oracle Support Note 1286796.1 `rp_filter` for multiple private interconnects and Linux Kernel 2.6.32+, which is available at the following URL:

<https://support.oracle.com/rs?type=doc&id=1286796.1>

Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle Database

Before installation, create operating system groups and users, and configure user environments.

[Creating Groups, Users and Paths for Oracle Grid Infrastructure](#) (page 6-1)

Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group, and create a software owner for Oracle Grid Infrastructure, and directories for Oracle home.

[Oracle Installations with Standard and Job Role Separation Groups and Users](#) (page 6-8)

A job role separation configuration of Oracle Database and Oracle ASM is a configuration with groups and users to provide separate groups for operating system authentication.

[Creating Operating System Privileges Groups](#) (page 6-11)

The following sections describe how to create operating system groups for Oracle Grid Infrastructure and Oracle Database:

[Creating Operating System Oracle Installation User Accounts](#) (page 6-14)

Before starting installation, create Oracle software owner user accounts, and configure their environments.

[Configuring Grid Infrastructure Software Owner User Environments](#) (page 6-22)

Understand the software owner user environments to configure before installing Oracle Grid Infrastructure.

[Enabling Intelligent Platform Management Interface \(IPMI\)](#) (page 6-28)

Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to computer hardware and firmware that system administrators can use to monitor system health and manage the system.

[Determining Root Script Execution Plan](#) (page 6-33)

During Oracle Grid Infrastructure installation, the installer requires you to run scripts with superuser (or `root`) privileges to complete a number of system configuration tasks.

6.1 Creating Groups, Users and Paths for Oracle Grid Infrastructure

Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group, and create a software owner for Oracle Grid Infrastructure, and directories for Oracle home.

Oracle software installations require an installation owner, an Oracle Inventory group, which is the primary group of all Oracle installation owners, and at least one group designated as a system privileges group. Review group and user options with your system administrator. If you have system administration privileges, then review the topics in this section and configure operating system groups and users as needed.

[Determining If an Oracle Inventory and Oracle Inventory Group Exist](#)

(page 6-2)

Determine if you have existing Oracle central inventory, and ensure that you use the same Oracle Inventory for all Oracle software installations. Also ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

[Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist](#)

(page 6-3)

If the `oraInst.loc` file does not exist, then create the Oracle Inventory group.

[About Oracle Installation Owner Accounts](#) (page 6-4)

Select or create an Oracle installation owner for your installation, depending on the group and user management plan you want to use for your installations.

[Restrictions for Oracle Software Installation Owners](#) (page 6-4)

Review the following restrictions for users created to own Oracle software.

[Identifying an Oracle Software Owner User Account](#) (page 6-5)

You must create at least one software owner user account the first time you install Oracle software on the system. Either use an existing Oracle software user account, or create an Oracle software owner user account for your installation.

[About the Oracle Base Directory for the Grid User](#) (page 6-6)

Review this information about creating the Oracle base directory on each cluster node.

[About the Oracle Home Directory for Oracle Grid Infrastructure Software](#)

(page 6-6)

Review this information about creating the Oracle home directory location on each cluster node.

[About Creating the Oracle Home and Oracle Base Directory](#) (page 6-7)

Create Grid home and Oracle base home directories on each cluster node.

6.1.1 Determining If an Oracle Inventory and Oracle Inventory Group Exist

Determine if you have existing Oracle central inventory, and ensure that you use the same Oracle Inventory for all Oracle software installations. Also ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

When you install Oracle software on the system for the first time, OUI creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group (by default, `oinstall`), and the path of the Oracle central inventory directory. If you have an existing Oracle central inventory, then ensure that you use the same Oracle

Inventory for all Oracle software installations, and ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

An `oraInst.loc` file contains lines in the following format, where `central_inventory_location` is the path to an existing Oracle central inventory, and `group` is the name of the operating system group whose members have permissions to write to the central inventory:

```
inventory_loc=central_inventory_location
inst_group=group
```

Use the `more` command to determine if you have an Oracle central inventory on your system. For example:

```
# more /etc/oraInst.loc

inventory_loc=/u01/app/oraInventory
inst_group=oinstall
```

Use the command `grep groupname /etc/group` to confirm that the group specified as the Oracle Inventory group still exists on the system. For example:

```
$ grep oinstall /etc/group
oinstall:x:54321:grid,oracle
```

Note: Do not put the `oraInventory` directory under the Oracle base directory for a new installation, because that can result in user permission errors for other installations.

6.1.2 Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist

If the `oraInst.loc` file does not exist, then create the Oracle Inventory group.

Members of the `OINSTALL` group are granted privileges to write to the Oracle central inventory (`oraInventory`), and other system privileges for Oracle installation owner users.

An Oracle installation owner should always have the group you want to have designated as the `OINSTALL` group (`oinstall`) as its primary group. Ensure that this group is available as the primary group for all planned Oracle software installation owners. By default, if an `oraInst.loc` file does not exist and an Oracle central inventory (`oraInventory`) is not identified, then the installer designates the primary group of the installation owner running the installation as the `OINSTALL` group.

The following example creates the `oraInventory` group `oinstall`, with the group ID number 54321.

```
# /usr/sbin/groupadd -g 54321 oinstall
```

Note:

For installations on Oracle Clusterware, group and user IDs must be identical on all nodes in the cluster. Ensure that the group and user IDs you want to use are available on each cluster member node, and confirm that the primary group for each Oracle Grid Infrastructure for a cluster installation owner has the same name and group ID.

6.1.3 About Oracle Installation Owner Accounts

Select or create an Oracle installation owner for your installation, depending on the group and user management plan you want to use for your installations.

You must create a software owner for your installation in the following circumstances:

- If an Oracle software owner user does not exist; for example, if this is the first installation of Oracle software on the system.
- If an Oracle software owner user exists, but you want to use a different operating system user, with different group membership, to separate Oracle Grid Infrastructure administrative privileges from Oracle Database administrative privileges.

In Oracle documentation, a user created to own only Oracle Grid Infrastructure software installations is called the Grid user (`grid`). This user owns both the Oracle Clusterware and Oracle Automatic Storage Management binaries. A user created to own either all Oracle installations, or one or more Oracle database installations, is called the Oracle user (`oracle`). You can have only one Oracle Grid Infrastructure installation owner, but you can have different Oracle users to own different installations.

Oracle software owners must have the Oracle Inventory group as their primary group, so that each Oracle software installation owner can write to the central inventory (`oraInventory`), and so that OCR and Oracle Clusterware resource permissions are set correctly. The database software owner must also have the OSDBA group and (if you create them) the OSOPER, OSBACKUPDBA, OSDGDBA, OSRACDBA, and OSKMDBA groups as secondary groups.

6.1.4 Restrictions for Oracle Software Installation Owners

Review the following restrictions for users created to own Oracle software.

- If you intend to use multiple Oracle software owners for different Oracle Database homes, then Oracle recommends that you create a separate software owner for Oracle Grid Infrastructure software (Oracle Clusterware and Oracle ASM), and use that owner to run the Oracle Grid Infrastructure installation.
- During installation, SSH must be set up between cluster member nodes. SSH can be set up automatically by Oracle Universal Installer (the installer). To enable SSH to be set up automatically, create Oracle installation owners without any `stty` commands in their profiles, and remove other security measures that are triggered during a login that generate messages to the terminal. These messages, mail checks, and other displays prevent Oracle software installation owner accounts from using the SSH configuration script that is built into the installer. If they are not disabled, then SSH must be configured manually before an installation can be run.
- If you plan to install Oracle Database or Oracle RAC, then Oracle recommends that you create separate users for the Oracle Grid Infrastructure and the Oracle Database installations. If you use one installation owner, then when you want to perform administration tasks, you must change the value for `$ORACLE_HOME` to the instance you want to administer (Oracle ASM, in the Oracle Grid Infrastructure home, or the database in the Oracle home), using command syntax such as the following example, where `/u01/app/12.2.0/grid` is the Oracle Grid Infrastructure home:

```
$ ORACLE_HOME=/u01/app/12.2.0/grid;
export ORACLE_HOME
```

- If you try to administer an Oracle home or Grid home instance using `sqlplus`, `lsnrctl`, or `asmcmd` commands while the environment variable `$ORACLE_HOME` is set to a different Oracle home or Grid home path, then you encounter errors. For example, when you start `SRVCTL` from a database home, `$ORACLE_HOME` should be set to that database home, or `SRVCTL` fails. The exception is when you are using `SRVCTL` in the Oracle Grid Infrastructure home. In that case, `$ORACLE_HOME` is ignored, and the Oracle home environment variable does not affect `SRVCTL` commands. In all other cases, you must change `$ORACLE_HOME` to the instance that you want to administer.
- To create separate Oracle software owners and separate operating system privileges groups for different Oracle software installations, note that each of these users must have the Oracle central inventory group (`oraInventory` group) as their primary group. Members of this group are granted the `OINSTALL` system privileges to write to the Oracle central inventory (`oraInventory`) directory, and are also granted permissions for various Oracle Clusterware resources, OCR keys, directories in the Oracle Clusterware home to which DBAs need write access, and other necessary privileges. Members of this group are also granted execute permissions to start and stop Clusterware infrastructure resources and databases. In Oracle documentation, this group is represented as `oinstall` in code examples.
- Each Oracle software owner must be a member of the same central inventory `oraInventory` group, and they must have this group as their primary group, so that all Oracle software installation owners share the same `OINSTALL` system privileges. Oracle recommends that you do not have more than one central inventory for Oracle installations. If an Oracle software owner has a different central inventory group, then you may corrupt the central inventory.

6.1.5 Identifying an Oracle Software Owner User Account

You must create at least one software owner user account the first time you install Oracle software on the system. Either use an existing Oracle software user account, or create an Oracle software owner user account for your installation.

To use an existing user account, obtain from you system administrator the name of an existing Oracle installation owner. Confirm that the existing owner is a member of the Oracle Inventory group.

For example if you know that the name of the Oracle Inventory group is `oinstall`, then an Oracle software owner should be listed as a member of `oinstall`:

```
$ grep "oinstall" /etc/group
oinstall:x:54321:grid,oracle
```

You can then use the `ID` command to verify that the Oracle installation owners you intend to use have the Oracle Inventory group as their primary group. For example:

```
$ id oracle
uid=54321(oracle) gid=54321(oinstall) groups=54321(oinstall),54322(dba),
54323(oper),54324(backupdba),54325(dgdba),54326(kmdba),54327(asmdba),54330(racdba)

$ id grid
uid=54331(grid) gid=54321(oinstall) groups=54321(oinstall),54322(dba),
54327(asmdba),54328(asmoper),54329(asmadmin),54330(racdba)
```

For Oracle Restart installations, to successfully install Oracle Database, ensure that the `grid` user is a member of the `racdba` group.

After you create operating system groups, create or modify Oracle user accounts in accordance with your operating system authentication planning.

6.1.6 About the Oracle Base Directory for the Grid User

Review this information about creating the Oracle base directory on each cluster node.

The Oracle base directory for the Oracle Grid Infrastructure installation is the location where diagnostic and administrative logs, and other logs associated with Oracle ASM and Oracle Clusterware are stored. For Oracle installations other than Oracle Grid Infrastructure for a cluster, it is also the location under which an Oracle home is placed.

However, in the case of an Oracle Grid Infrastructure installation, you must create a *different* path, so that the path for Oracle bases remains available for other Oracle installations.

For OUI to recognize the Oracle base path, it must be in the form `u[00-99][00-99]/app`, and it must be writable by any member of the `oraInventory` (`oinstall`) group. The OFA path for the Oracle base is `u[00-99][00-99]/app/user`, where *user* is the name of the software installation owner. For example:

```
/u01/app/grid
```

6.1.7 About the Oracle Home Directory for Oracle Grid Infrastructure Software

Review this information about creating the Oracle home directory location on each cluster node.

The Oracle home for Oracle Grid Infrastructure software (Grid home) should be located in a path that is different from the Oracle home directory paths for any other Oracle software. The Optimal Flexible Architecture guideline for a Grid home is to create a path in the form `/p/m/v/u`, where `p` is a string constant, `m` is a unique fixed-length key (typically a two-digit number), `v` is the version of the software, and `u` is the installation owner of the Oracle Grid Infrastructure software (`grid` user). During Oracle Grid Infrastructure for a cluster installation, the path of the Grid home is changed to the root user, so any other users are unable to read, write, or execute commands in that path. For example, to create a Grid home in the standard mount point path format `u[00-99][00-99]/app/release/grid`, where *release* is the release number of the Oracle Grid Infrastructure software, create the following path:

```
/u01/app/12.2.0/grid
```

During installation, ownership of the entire path to the Grid home is changed to `root` (`/u01`, `/u01/app`, `/u01/app/12.2.0`, `/u01/app/12.2.0/grid`). If you do not create a unique path to the Grid home, then after the Grid install, you can encounter permission errors for other installations, including any existing installations under the same path. To avoid placing the application directory in the mount point under root ownership, you can create and select paths such as the following for the Grid home:

```
/u01/12.2.0/grid
```

Caution:

For Oracle Grid Infrastructure for a cluster installations, note the following restrictions for the Oracle Grid Infrastructure binary home (Grid home directory for Oracle Grid Infrastructure):

- It must not be placed under one of the Oracle base directories, including the Oracle base directory of the Oracle Grid Infrastructure installation owner.
- It must not be placed in the home directory of an installation owner. These requirements are specific to Oracle Grid Infrastructure for a cluster installations.

Oracle Grid Infrastructure for a standalone server (Oracle Restart) can be installed under the Oracle base for the Oracle Database installation.

6.1.8 About Creating the Oracle Home and Oracle Base Directory

Create Grid home and Oracle base home directories on each cluster node.

Oracle recommends that you create Oracle Grid Infrastructure Grid home and Oracle base homes manually, particularly if you have separate Oracle Grid Infrastructure for a cluster and Oracle Database software owners, so that you can separate log files for the Oracle Grid Infrastructure installation owner in a separate Oracle base, and prevent accidental placement of the Grid home under an Oracle base path. For example:

```
# mkdir -p /u01/app/12.2.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R grid:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/
```

Note:

Placing Oracle Grid Infrastructure for a cluster binaries on a cluster file system is not supported.

If you plan to install an Oracle RAC home on a shared OCFS2 location, then you must upgrade OCFS2 to at least version 1.4.1, which supports shared writable mmapS.

Oracle recommends that you install Oracle Grid Infrastructure locally, on each cluster member node. Using a shared Grid home prevents rolling upgrades, and creates a single point of failure for the cluster.

Related Topics:

[Optimal Flexible Architecture](#) (page C-1)

Oracle Optimal Flexible Architecture (OFA) rules are a set of configuration guidelines created to ensure well-organized Oracle installations, which simplifies administration, support and maintenance.

6.2 Oracle Installations with Standard and Job Role Separation Groups and Users

A job role separation configuration of Oracle Database and Oracle ASM is a configuration with groups and users to provide separate groups for operating system authentication.

Review the following sections to understand more about a Job Role Separation deployment:

[About Oracle Installations with Job Role Separation](#) (page 6-8)

Job role separation requires that you create different operating system groups for each set of system privileges that you grant through operating system authorization.

[Standard Oracle Database Groups for Database Administrators](#) (page 6-9)

Oracle Database has two standard administration groups: OSDBA, which is required, and OSOPER, which is optional.

[Extended Oracle Database Groups for Job Role Separation](#) (page 6-9)

Oracle Database 12c Release 1 (12.1) and later releases provide an extended set of database groups to grant task-specific system privileges for database administration.

[Creating an ASMSNMP User](#) (page 6-10)

The ASMSNMP user is an Oracle ASM user with privileges to monitor Oracle ASM instances. You are prompted to provide a password for this user during installation.

[Oracle Automatic Storage Management Groups for Job Role Separation](#) (page 6-11)

Oracle Grid Infrastructure operating system groups provide their members task-specific system privileges to access and to administer Oracle Automatic Storage Management.

6.2.1 About Oracle Installations with Job Role Separation

Job role separation requires that you create different operating system groups for each set of system privileges that you grant through operating system authorization.

With Oracle Grid Infrastructure job role separation, Oracle ASM has separate operating system groups that provide operating system authorization for Oracle ASM system privileges for storage tier administration. This operating system authorization is separated from Oracle Database operating system authorization. In addition, the Oracle Grid Infrastructure installation owner provides operating system user authorization for modifications to Oracle Grid Infrastructure binaries.

With Oracle Database job role separation, each Oracle Database installation has separate operating system groups to provide authorization for system privileges on that Oracle Database. Multiple databases can, therefore, be installed on the cluster without sharing operating system authorization for system privileges. In addition, each Oracle software installation is owned by a separate installation owner, to provide operating system user authorization for modifications to Oracle Database binaries. Note that any Oracle software owner can start and stop all databases and shared Oracle Grid Infrastructure resources such as Oracle ASM or Virtual IP (VIP). Job role separation configuration enables database security, and does not restrict user roles in starting and stopping various Oracle Clusterware resources.

You can choose to create one administrative user and one group for operating system authentication for all system privileges on the storage and database tiers. For example, you can designate the `oracle` user to be the installation owner for all Oracle software, and designate `oinstall` to be the group whose members are granted all system privileges for Oracle Clusterware; all system privileges for Oracle ASM; all system privileges for all Oracle Databases on the servers; and all `OINSTALL` system privileges for installation owners. This group must also be the Oracle Inventory group.

If you do not want to use role allocation groups, then Oracle strongly recommends that you use at least two groups:

- A system privileges group whose members are granted administrative system privileges, including `OSDBA`, `OSASM`, and other system privileges groups.
- An installation owner group (the `oraInventory` group) whose members are granted Oracle installation owner system privileges (the `OINSTALL` system privilege).

Note: To configure users for installation that are on a network directory service such as Network Information Services (NIS), refer to your directory service documentation.

Related Topics:

Oracle Database Administrator's Guide

Oracle Automatic Storage Management Administrator's Guide

6.2.2 Standard Oracle Database Groups for Database Administrators

Oracle Database has two standard administration groups: `OSDBA`, which is required, and `OSOPER`, which is optional.

- The `OSDBA` group (typically, `dba`)

You must create this group the first time you install Oracle Database software on the system. This group identifies operating system user accounts that have database administrative privileges (the `SYSDBA` privilege).

If you do not create separate `OSDBA`, `OSOPER`, and `OSASM` groups for the Oracle ASM instance, then operating system user accounts that have the `SYSOPER` and `SYSASM` privileges must be members of this group. The name used for this group in Oracle code examples is `dba`. If you do not designate a separate group as the `OSASM` group, then the `OSDBA` group you define is also by default the `OSASM` group.

- The `OSOPER` group for Oracle Database (typically, `oper`)

`OSOPER` grants the `OPERATOR` privilege to start up and shut down the database (the `SYSOPER` privilege). By default, members of the `OSDBA` group have all privileges granted by the `SYSOPER` privilege.

6.2.3 Extended Oracle Database Groups for Job Role Separation

Oracle Database 12c Release 1 (12.1) and later releases provide an extended set of database groups to grant task-specific system privileges for database administration.

The extended set of Oracle Database system privileges groups are task-specific and less privileged than the OSDBA/SYSDBA system privileges. They are designed to provide privileges to carry out everyday database operations. Users granted these system privileges are also authorized through operating system group membership.

You do not have to create these specific group names, but during interactive and silent installation, you must assign operating system groups whose members are granted access to these system privileges. You can assign the same group to provide authorization for these privileges, but Oracle recommends that you provide a unique group to designate each privilege.

The subset of OSDBA job role separation privileges and groups consist of the following:

- OSBACKUPDBA group for Oracle Database (typically, `backupdba`)
Create this group if you want a separate group of operating system users to have a limited set of database backup and recovery related administrative privileges (the SYSBACKUP privilege).
- OSDGDBA group for Oracle Data Guard (typically, `dgdba`)
Create this group if you want a separate group of operating system users to have a limited set of privileges to administer and monitor Oracle Data Guard (the SYSDG privilege). To use this privilege, add the Oracle Database installation owners as members of this group.
- The OSKMDBA group for encryption key management (typically, `kmdba`)
Create this group if you want a separate group of operating system users to have a limited set of privileges for encryption key management such as Oracle Wallet Manager management (the SYSKM privilege). To use this privilege, add the Oracle Database installation owners as members of this group.
- The OSRACDBA group for Oracle Real Application Clusters Administration (typically, `racdba`)
Create this group if you want a separate group of operating system users to have a limited set of Oracle Real Application Clusters (RAC) administrative privileges (the SYSRAC privilege). To use this privilege:
 - Add the Oracle Database installation owners as members of this group.
 - For Oracle Restart configurations, if you have a separate Oracle Grid Infrastructure installation owner user (`grid`), then you must also add the `grid` user as a member of the OSRACDBA group of the database to enable Oracle Grid Infrastructure components to connect to the database.

Related Topics:

Oracle Database Administrator's Guide

Oracle Database Security Guide

6.2.4 Creating an ASMSNMP User

The ASMSNMP user is an Oracle ASM user with privileges to monitor Oracle ASM instances. You are prompted to provide a password for this user during installation.

In addition to the OSASM group, whose members are granted the SYSASM system privilege to administer Oracle ASM, Oracle recommends that you create a less

privileged user, ASMSNMP, and grant that user SYSDBA privileges to monitor the Oracle ASM instance. Oracle Enterprise Manager uses the ASMSNMP user to monitor Oracle ASM status.

During installation, you are prompted to provide a password for the ASMSNMP user. You can create an operating system authenticated user, or you can create an Oracle Database user called `asmsnmp`. In either case, grant the user SYSDBA privileges.

6.2.5 Oracle Automatic Storage Management Groups for Job Role Separation

Oracle Grid Infrastructure operating system groups provide their members task-specific system privileges to access and to administer Oracle Automatic Storage Management.

- The OSASM group for Oracle ASM Administration (typically, `asmadmin`)

Create this group as a separate group to separate administration privileges groups for Oracle ASM and Oracle Database administrators. Members of this group are granted the SYSASM system privileges to administer Oracle ASM. In Oracle documentation, the operating system group whose members are granted privileges is called the OSASM group, and in code examples, where there is a group specifically created to grant this privilege, it is referred to as `asmadmin`.

Oracle ASM can support multiple databases. If you have multiple databases on your system, and use multiple OSDBA groups so that you can provide separate SYSDBA privileges for each database, then you should create a group whose members are granted the OSASM/SYSASM administrative privileges, and create a grid infrastructure user (`grid`) that does not own a database installation, so that you separate Oracle Grid Infrastructure SYSASM administrative privileges from a database administrative privileges group.

Members of the OSASM group can use SQL to connect to an Oracle ASM instance as SYSASM using operating system authentication. The SYSASM privileges permit mounting and dismounting disk groups, and other storage administration tasks. SYSASM privileges provide no access privileges on an RDBMS instance.

If you do not designate a separate group as the OSASM group, but you do define an OSDBA group for database administration, then by default the OSDBA group you define is also defined as the OSASM group.

- The OSOPER group for Oracle ASM (typically, `asmoper`)

This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of Oracle instance administrative privileges (the SYSOPER for ASM privilege), including starting up and stopping the Oracle ASM instance. By default, members of the OSASM group also have all privileges granted by the SYSOPER for ASM privilege.

6.3 Creating Operating System Privileges Groups

The following sections describe how to create operating system groups for Oracle Grid Infrastructure and Oracle Database:

[Creating the OSASM Group](#) (page 6-12)

If the OSASM group does not exist, or if you require a new OSASM group, then create it.

[Creating the OSDBA for ASM Group](#) (page 6-12)

You must designate a group as the OSDBA for ASM (asmdba) group during installation. Members of this group are granted access privileges to Oracle Automatic Storage Management.

[Creating the OSOPER for ASM Group](#) (page 6-13)

You can choose to designate a group as the OSOPER for ASM group (asmoper) during installation. Members of this group are granted startup and shutdown privileges to Oracle Automatic Storage Management.

[Creating the OSDBA Group for Database Installations](#) (page 6-13)

Each Oracle Database requires an operating system group to be designated as the OSDBA group. Members of this group are granted the SYSDBA system privileges to administer the database.

[Creating an OSOPER Group for Database Installations](#) (page 6-13)

Create an OSOPER group only if you want to identify a group of operating system users with a limited set of database administrative privileges (SYSOPER operator privileges).

[Creating the OSBACKUPDBA Group for Database Installations](#) (page 6-14)

You must designate a group as the OSBACKUPDBA group during installation. Members of this group are granted the SYSBACKUP privileges to perform backup and recovery operations using RMAN or SQL*Plus.

[Creating the OSDGDBA Group for Database Installations](#) (page 6-14)

You must designate a group as the OSDGDBA group during installation. Members of this group are granted the SYSDG privileges to perform Data Guard operations.

[Creating the OSKMDBA Group for Database Installations](#) (page 6-14)

You must designate a group as the OSKMDBA group during installation. Members of this group are granted the SYSKM privileges to perform Transparent Data Encryption keystore operations.

[Creating the OSRACDBA Group for Database Installations](#) (page 6-14)

You must designate a group as the OSRACDBA group during database installation. Members of this group are granted the SYSRAC privileges to perform day-to-day administration of Oracle databases on an Oracle RAC cluster.

6.3.1 Creating the OSASM Group

If the OSASM group does not exist, or if you require a new OSASM group, then create it.

Use the group name asmadmin unless a group with that name already exists. For example:

```
# groupadd -g 54329 asmadmin
```

6.3.2 Creating the OSDBA for ASM Group

You must designate a group as the OSDBA for ASM (asmdba) group during installation. Members of this group are granted access privileges to Oracle Automatic Storage Management.

Create an OSDBA for ASM group using the group name `asmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54327 asmdba
```

6.3.3 Creating the OSOPER for ASM Group

You can choose to designate a group as the OSOPER for ASM group (`asmoper`) during installation. Members of this group are granted startup and shutdown privileges to Oracle Automatic Storage Management.

If you want to create an OSOPER for ASM group, use the group name `asmoper` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54328 asmoper
```

6.3.4 Creating the OSDBA Group for Database Installations

Each Oracle Database requires an operating system group to be designated as the OSDBA group. Members of this group are granted the SYSDBA system privileges to administer the database.

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSDBA group exists, but you want to give a different group of operating system users database administrative privileges for a new Oracle Database installation

Create the OSDBA group using the group name `dba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54322 dba
```

6.3.5 Creating an OSOPER Group for Database Installations

Create an OSOPER group only if you want to identify a group of operating system users with a limited set of database administrative privileges (SYSOPER operator privileges).

For most installations, it is sufficient to create only the OSDBA group. However, to use an OSOPER group, create it in the following circumstances:

- If an OSOPER group does not exist; for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of operating system users database operator privileges in a new Oracle installation

If the OSOPER group does not exist, or if you require a new OSOPER group, then create it. Use the group name `oper` unless a group with that name already exists. For example:

```
# groupadd -g 54323 oper
```

6.3.6 Creating the OSBACKUPDBA Group for Database Installations

You must designate a group as the OSBACKUPDBA group during installation. Members of this group are granted the SYSBACKUP privileges to perform backup and recovery operations using RMAN or SQL*Plus.

Create the OSBACKUPDBA group using the group name `backupdba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54324 backupdba
```

6.3.7 Creating the OSDGDBA Group for Database Installations

You must designate a group as the OSDGDBA group during installation. Members of this group are granted the SYSDG privileges to perform Data Guard operations.

Create the OSDGDBA group using the group name `dgdba`, unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54325 dgdba
```

6.3.8 Creating the OSKMDBA Group for Database Installations

You must designate a group as the OSKMDBA group during installation. Members of this group are granted the SYSKM privileges to perform Transparent Data Encryption keystore operations.

If you want a separate group for Transparent Data Encryption, then create the OSKMDBA group using the group name `kmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54326 kmdba
```

6.3.9 Creating the OSRACDBA Group for Database Installations

You must designate a group as the OSRACDBA group during database installation. Members of this group are granted the SYSRAC privileges to perform day-to-day administration of Oracle databases on an Oracle RAC cluster.

Create the OSRACDBA group using the group name `racdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 54330 racdba
```

6.4 Creating Operating System Oracle Installation User Accounts

Before starting installation, create Oracle software owner user accounts, and configure their environments.

Oracle software owner user accounts require resource settings and other environment configuration. To protect against accidents, Oracle recommends that you create one software installation owner account for each Oracle software program you install.

[Creating an Oracle Software Owner User](#) (page 6-15)

If the Oracle software owner user (`oracle` or `grid`) does not exist, or if you require a new Oracle software owner user, then create it as described in this section.

[Modifying Oracle Owner User Groups](#) (page 6-15)

If you have created an Oracle software installation owner account, but it is not a member of the groups you want to designate as the OSDBA, OSOPER, OSDBA for ASM, ASMADMIN, or other system privileges group, then modify the group settings for that user before installation.

[Identifying Existing User and Group IDs](#) (page 6-16)

To create identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, and then create the user and groups with the same name and ID on the other cluster nodes.

[Creating Identical Database Users and Groups on Other Cluster Nodes](#) (page 6-16)

Oracle software owner users and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes.

[Example of Creating Minimal Groups, Users, and Paths](#) (page 6-17)

You can create a minimal operating system authentication configuration as described in this example.

[Example of Creating Role-allocated Groups, Users, and Paths](#) (page 6-19)

Understand this example of how to create role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

6.4.1 Creating an Oracle Software Owner User

If the Oracle software owner user (`oracle` or `grid`) does not exist, or if you require a new Oracle software owner user, then create it as described in this section.

The following example shows how to create the user `oracle` with the user ID 54321; with the primary group `oinstall`; and with secondary groups `dba`, `asmdba`, `backupdba`, `dgdba`, `kmdba`, and `racdba`:

```
# /usr/sbin/useradd -u 54321 -g oinstall -G dba,asmdba,backupdba,dgdba,kmdba,racdba oracle
```

You must note the user ID number for installation users, because you need it during preinstallation.

For Oracle Grid Infrastructure installations, user IDs and group IDs must be identical on all candidate nodes.

6.4.2 Modifying Oracle Owner User Groups

If you have created an Oracle software installation owner account, but it is not a member of the groups you want to designate as the OSDBA, OSOPER, OSDBA for ASM, ASMADMIN, or other system privileges group, then modify the group settings for that user before installation.

Warning:

Each Oracle software owner must be a member of the same central inventory group. Do not modify the primary group of an existing Oracle software owner account, or designate different groups as the OINSTALL group. If Oracle software owner accounts have different groups as their primary group, then you can corrupt the central inventory.

During installation, the user that is installing the software should have the OINSTALL group as its primary group, and it must be a member of the operating system groups appropriate for your installation. For example:

```
# /usr/sbin/usermod -g oinstall -G dba,asmdba,backupdba,dgdba,kmdba,racdba[,oper]
oracle
```

6.4.3 Identifying Existing User and Group IDs

To create identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, and then create the user and groups with the same name and ID on the other cluster nodes.

1. Enter a command similar to the following (in this case, to determine a user ID for the `oracle` user):

```
# id oracle
```

The output from this command is similar to the following:

```
uid=54321(oracle) gid=54421(oinstall) groups=54322(dba),54323(oper),54327(asmdba)
```

2. From the output, identify the user ID (`uid`) for the user and the group identities (`gids`) for the groups to which it belongs.

Ensure that these ID numbers are identical on each node of the cluster. The user's primary group is listed after `gid`. Secondary groups are listed after `groups`.

6.4.4 Creating Identical Database Users and Groups on Other Cluster Nodes

Oracle software owner users and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes.

To create users and groups on the other cluster nodes, repeat the following procedure on each node:

You must complete the following procedures only if you are using local users and groups. If you are using users and groups defined in a directory service such as NIS, then they are already identical on each cluster node.

1. Log in to the node as `root`.
2. Enter commands similar to the following to create the `asmadmin`, `asmdba`, `backupdba`, `dgdba`, `kmdba`, `asmoper` and `oper` groups, and if not configured by the Oracle Preinstallation RPM or prior installations, then the `oinstall` and `dba` groups.

Use the `-g` option to specify the correct group ID for each group.

```
# groupadd -g 54421 oinstall
# groupadd -g 54322 dba
```

```
# groupadd -g 54323 oper
# groupadd -g 54324 backupdba
# groupadd -g 54325 dgdba
# groupadd -g 54326 kmdba
# groupadd -g 54327 asmdba
# groupadd -g 54328 asmoper
# groupadd -g 54329 asmadmin
# groupadd -g 54330 racdba
```

Note: You are not required to use the UIDs and GIDs in this example. If a group already exists, then use the `groupmod` command to modify it if necessary. If you cannot use the same group ID for a particular group on a node, then view the `/etc/group` file on all nodes to identify a group ID that is available on every node. You must then change the group ID on all nodes to the same group ID.

- To create the oracle or Oracle Grid Infrastructure (`grid`) user, enter a command similar to the following:

```
# useradd -u 54322 -g oinstall -G asmadmin,asmdba grid
```

- The `-u` option specifies the user ID, which must be the user ID that you identified earlier.
- The `-g` option specifies the primary group for the Grid user, which must be the Oracle Inventory group (OINSTALL), which grants the OINSTALL system privileges. In this example, the OINSTALL group is `oinstall`.
- The `-G` option specifies the secondary groups. The Grid user must be a member of the OSASM group (`asmadmin`) and the OSDBA for ASM group (`asmdba`).

Note: If the user already exists, then use the `usermod` command to modify it if necessary. If you cannot use the same user ID for the user on every node, then view the `/etc/passwd` file on all nodes to identify a user ID that is available on every node. You must then specify that ID for the user on all of the nodes.

- Set the password of the user.

For example:

```
# passwd oracle
```

- Complete user environment configuration tasks for each user.

6.4.5 Example of Creating Minimal Groups, Users, and Paths

You can create a minimal operating system authentication configuration as described in this example.

This configuration example shows the following:

- Creation of the Oracle Inventory group (`oinstall`)
- Creation of a single group (`dba`) as the only system privileges group to assign for all Oracle Grid Infrastructure, Oracle ASM, and Oracle Database system privileges

- Creation of the Oracle Grid Infrastructure software owner (`grid`), and one Oracle Database owner (`oracle`) with correct group memberships
- Creation and configuration of an Oracle base path compliant with OFA structure with correct permissions

Enter the following commands to create a minimal operating system authentication configuration:

```
# groupadd -g 54421 oinstall
# groupadd -g 54422 dba
# useradd -u 54321 -g oinstall -G dba oracle
# useradd -u 54322 -g oinstall -G dba grid
# mkdir -p /u01/app/12.2.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R grid:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/
```

After running these commands, you have the following groups and users:

- An Oracle central inventory group, or oraInventory group (`oinstall`). Members who have the central inventory group as their primary group, are granted the OINSTALL permission to write to the oraInventory directory.
- One system privileges group, `dba`, for Oracle Grid Infrastructure, Oracle ASM and Oracle Database system privileges. Members who have the `dba` group as their primary or secondary group are granted operating system authentication for OSASM/SYSASM, OSDBA/SYSDBA, OSOPER/SYSOPER, OSBACKUPDBA/SYSBACKUP, OSDGDBA/SYSDG, OSKMDBA/SYSKM, OSDBA for ASM/SYSDBA for ASM, and OSOPER for ASM/SYSOPER for Oracle ASM to administer Oracle Clusterware, Oracle ASM, and Oracle Database, and are granted SYSASM and OSOPER for Oracle ASM access to the Oracle ASM storage.
- An Oracle Grid Infrastructure for a cluster owner, or Grid user (`grid`), with the oraInventory group (`oinstall`) as its primary group, and with the OSASM group (`dba`) as the secondary group, with its Oracle base directory `/u01/app/grid`.
- An Oracle Database owner (`oracle`) with the oraInventory group (`oinstall`) as its primary group, and the OSDBA group (`dba`) as its secondary group, with its Oracle base directory `/u01/app/oracle`.
- `/u01/app` owned by `grid:oinstall` with 775 permissions before installation, and by root after the `root.sh` script is run during installation. This ownership and permissions enables OUI to create the Oracle Inventory directory, in the path `/u01/app/oraInventory`.
- `/u01` owned by `grid:oinstall` before installation, and by root after the `root.sh` script is run during installation.
- `/u01/app/12.2.0/grid` owned by `grid:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.
- `/u01/app/grid` owned by `grid:oinstall` with 775 permissions. These permissions are required for installation, and are changed during the installation process.

- /u01/app/oracle owned by oracle:oinstall with 775 permissions.

Note: You can use one installation owner for both Oracle Grid Infrastructure and any other Oracle installations. However, Oracle recommends that you use separate installation owner accounts for each Oracle software installation.

6.4.6 Example of Creating Role-allocated Groups, Users, and Paths

Understand this example of how to create role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

This example illustrates the following scenario:

- An Oracle Grid Infrastructure installation
- Two separate Oracle Database installations planned for the cluster, DB1 and DB2
- Separate installation owners for Oracle Grid Infrastructure, and for each Oracle Database
- Full role allocation of system privileges for Oracle ASM, and for each Oracle Database
- Oracle Database owner oracle1 granted the right to start up and shut down the Oracle ASM instance

Create groups and users for a role-allocated configuration for this scenario using the following commands:

```
# groupadd -g 54321 oinstall
# groupadd -g 54322 dba1
# groupadd -g 54332 dba2
# groupadd -g 54323 oper1
# groupadd -g 54333 oper2
# groupadd -g 54324 backupdba1
# groupadd -g 54334 backupdba2
# groupadd -g 54325 dgdba1
# groupadd -g 54335 dgdba2
# groupadd -g 54326 kmdba1
# groupadd -g 54336 kmdba2
# groupadd -g 54327 asmdba
# groupadd -g 54328 asmoper
# groupadd -g 54329 asmadmin
# groupadd -g 54330 racdba1
# groupadd -g 54340 racdba2
# useradd -u 54322 -g oinstall -G asmadmin,asmdba,racdba1,racdba2 grid
# useradd -u 54321 -g oinstall -G
dba1,backupdba1,dgdba1,kmdba1,asmdba,racdba1,asmoper oracle1
# useradd -u 54323 -g oinstall -G dba2,backupdba2,dgdba2,kmdba2,asmdba,racdba2
oracle2
# mkdir -p /u01/app/12.2.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle1
# mkdir -p /u01/app/oracle2
# chown -R grid:oinstall /u01
# chmod -R 775 /u01/
# chown oracle1:oinstall /u01/app/oracle1
# chown oracle2:oinstall /u01/app/oracle2
```

After running these commands, you have a set of administrative privileges groups and users for Oracle Grid Infrastructure, and for two separate Oracle databases (DB1 and DB2):

Example 6-1 Oracle Grid Infrastructure Groups and Users Example

The command creates the following Oracle Grid Infrastructure groups and users:

- An Oracle central inventory group, or `oraInventory` group (`oinstall`), whose members that have this group as their primary group. Members of this group are granted the `OINSTALL` system privileges, which grants permissions to write to the `oraInventory` directory, and other associated install binary privileges.
- An OSASM group (`asmadmin`), associated with Oracle Grid Infrastructure during installation, whose members are granted the `SYSASM` privileges to administer Oracle ASM.
- An OSDBA for ASM group (`asmdba`), associated with Oracle Grid Infrastructure storage during installation. Its members include `grid` and any database installation owners, such as `oracle1` and `oracle2`, who are granted access to Oracle ASM. Any additional installation owners that use Oracle ASM for storage must also be made members of this group.
- An OSOPER for ASM group for Oracle ASM (`asmoper`), associated with Oracle Grid Infrastructure during installation. Members of `asmoper` group are granted limited Oracle ASM administrator privileges, including the permissions to start and stop the Oracle ASM instance.
- An Oracle Grid Infrastructure installation owner (`grid`), with the `oraInventory` group (`oinstall`) as its primary group, and with the OSASM (`asmadmin`) group and the OSDBA for ASM (`asmdba`) group as secondary groups.
- `/u01/app/oraInventory`. The central inventory of Oracle installations on the cluster. This path remains owned by `grid:oinstall`, to enable other Oracle software owners to write to the central inventory.
- An OFA-compliant mount point `/u01` owned by `grid:oinstall` before installation, so that Oracle Universal Installer can write to that path.
- An Oracle base for the `grid` installation owner `/u01/app/grid` owned by `grid:oinstall` with 775 permissions, and changed during the installation process to 755 permissions.
- A Grid home `/u01/app/12.1.0/grid` owned by `grid:oinstall` with 775 (`drwxdrwxr-x`) permissions. These permissions are required for installation, and are changed during the installation process to `root:oinstall` with 755 permissions (`drwxr-xr-x`).

Example 6-2 Oracle Database DB1 Groups and Users Example

The command creates the following Oracle Database (DB1) groups and users:

- An Oracle Database software owner (`oracle1`), which owns the Oracle Database binaries for DB1. The `oracle1` user has the `oraInventory` group as its primary group, and the OSDBA group for its database (`dba1`) and the OSDBA for ASM group for Oracle Grid Infrastructure (`asmdba`) as secondary groups. In addition, the `oracle1` user is a member of `asmoper`, granting that user privileges to start up and shut down Oracle ASM.

- An OSDBA group (`dba1`). During installation, you identify the group `dba1` as the OSDBA group for the database installed by the user `oracle1`. Members of `dba1` are granted the SYSDBA privileges for the Oracle Database DB1. Users who connect as SYSDBA are identified as user SYS on DB1.
- An OSBACKUPDBA group (`backupdba1`). During installation, you identify the group `backupdba1` as the OSDBA group for the database installed by the user `oracle1`. Members of `backupdba1` are granted the SYSBACKUP privileges for the database installed by the user `oracle1` to back up the database.
- An OSDGDBA group (`dgdba1`). During installation, you identify the group `dgdba1` as the OSDGDBA group for the database installed by the user `oracle1`. Members of `dgdba1` are granted the SYSDG privileges to administer Oracle Data Guard for the database installed by the user `oracle1`.
- An OSKMDBA group (`kmdba1`). During installation, you identify the group `kmdba1` as the OSKMDBA group for the database installed by the user `oracle1`. Members of `kmdba1` are granted the SYSKM privileges to administer encryption keys for the database installed by the user `oracle1`.
- An OSOPER group (`oper1`). During installation, you identify the group `oper1` as the OSOPER group for the database installed by the user `oracle1`. Members of `oper1` are granted the SYSOPER privileges (a limited set of the SYSDBA privileges), including the right to start up and shut down the DB1 database. Users who connect as OSOPER privileges are identified as user PUBLIC on DB1.
- An Oracle base `/u01/app/oracle1` owned by `oracle1:oinstall` with 775 permissions. The user `oracle1` has permissions to install software in this directory, but in no other directory in the `/u01/app` path.

Example 6-3 Oracle Database DB2 Groups and Users Example

The command creates the following Oracle Database (DB2) groups and users:

- An Oracle Database software owner (`oracle2`), which owns the Oracle Database binaries for DB2. The `oracle2` user has the `oraInventory` group as its primary group, and the OSDBA group for its database (`dba2`) and the OSDBA for ASM group for Oracle Grid Infrastructure (`asmdba`) as secondary groups. However, the `oracle2` user is not a member of the `asmoper` group, so `oracle2` cannot shut down or start up Oracle ASM.
- An OSDBA group (`dba2`). During installation, you identify the group `dba2` as the OSDBA group for the database installed by the user `oracle2`. Members of `dba2` are granted the SYSDBA privileges for the Oracle Database DB2. Users who connect as SYSDBA are identified as user SYS on DB2.
- An OSBACKUPDBA group (`backupdba2`). During installation, you identify the group `backupdba2` as the OSDBA group for the database installed by the user `oracle2`. Members of `backupdba2` are granted the SYSBACKUP privileges for the database installed by the user `oracle2` to back up the database.
- An OSDGDBA group (`dgdba2`). During installation, you identify the group `dgdba2` as the OSDGDBA group for the database installed by the user `oracle2`. Members of `dgdba2` are granted the SYSDG privileges to administer Oracle Data Guard for the database installed by the user `oracle2`.
- An OSKMDBA group (`kmdba2`). During installation, you identify the group `kmdba2` as the OSKMDBA group for the database installed by the user `oracle2`.

Members of `kmdba2` are granted the `SYSKM` privileges to administer encryption keys for the database installed by the user `oracle2`.

- An OSOPER group (`oper2`). During installation, you identify the group `oper2` as the OSOPER group for the database installed by the user `oracle2`. Members of `oper2` are granted the `SYSOPER` privileges (a limited set of the `SYSDBA` privileges), including the right to start up and shut down the DB2 database. Users who connect as OSOPER privileges are identified as user `PUBLIC` on DB2.
- An Oracle base `/u01/app/oracle2` owned by `oracle1:oinstall` with 775 permissions. The user `oracle2` has permissions to install software in this directory, but in no other directory in the `/u01/app` path.

6.5 Configuring Grid Infrastructure Software Owner User Environments

Understand the software owner user environments to configure before installing Oracle Grid Infrastructure.

You run the installer software with the Oracle Grid Infrastructure installation owner user account (`oracle` or `grid`). However, before you start the installer, you must configure the environment of the installation owner user account. If needed, you must also create other required Oracle software owners.

[Environment Requirements for Oracle Software Owners](#) (page 6-22)

You must make the following changes to configure Oracle software owner environments:

[Procedure for Configuring Oracle Software Owner Environments](#) (page 6-23)

Configure each Oracle installation owner user account environment:

[Checking Resource Limits for Oracle Software Installation Users](#) (page 6-25)

For each installation software owner user account, check the resource limits for installation.

[Setting Remote Display and X11 Forwarding Configuration](#) (page 6-26)

If you are on a remote terminal, and the local system has only one visual (which is typical), then use the following syntax to set your user account `DISPLAY` environment variable:

[Preventing Installation Errors Caused by Terminal Output Commands](#)

(page 6-27)

During an Oracle Grid Infrastructure installation, OUI uses SSH to run commands and copy files to the other nodes. During the installation, hidden files on the system (for example, `.bashrc` or `.cshrc`) can cause `makefile` and other installation errors if they contain terminal output commands.

6.5.1 Environment Requirements for Oracle Software Owners

You must make the following changes to configure Oracle software owner environments:

- Set the installation software owner user (`grid`, `oracle`) default file mode creation mask (`umask`) to 022 in the shell startup file. Setting the mask to 022 ensures that the user performing the software installation creates files with 644 permissions.

- Set ulimit settings for file descriptors and processes for the installation software owner (`grid`, `oracle`).
- Set the `DISPLAY` environment variable in preparation for running an Oracle Universal Installer (OUI) installation.

Caution:

If you have existing Oracle installations that you installed with the user ID that is your Oracle Grid Infrastructure software owner, then unset all Oracle environment variable settings for that user.

6.5.2 Procedure for Configuring Oracle Software Owner Environments

Configure each Oracle installation owner user account environment:

1. Start an X terminal session (`xterm`) on the server where you are running the installation.
2. Enter the following command to ensure that X Window applications can display on this system, where `hostname` is the fully qualified name of the local host from which you are accessing the server:

```
$ xhost + hostname
```

3. If you are not logged in as the software owner user, then switch to the software owner user you are configuring. For example, with the user `grid`:

```
$ su - grid
```

On systems where you cannot run `su` commands, use `sudo` instead:

```
$ sudo -u grid -s
```

4. To determine the default shell for the user, enter the following command:

```
$ echo $SHELL
```

5. Open the user's shell startup file in any text editor:

- Bash shell (`bash`):

```
$ vi .bash_profile
```

- Bourne shell (`sh`) or Korn shell (`ksh`):

```
$ vi .profile
```

- C shell (`csh` or `tcsh`):

```
% vi .login
```

6. Enter or edit the following line, specifying a value of `022` for the default file mode creation mask:

```
umask 022
```

7. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, then remove these lines from the file.

8. Save the file, and exit from the text editor.
9. To run the shell startup script, enter one of the following commands:

- Bash shell:

```
$ . ./bash_profile
```
- Bourne, Bash, or Korn shell:

```
$ . ./profile
```
- C shell:

```
% source ./login
```

10. Use the following command to check the PATH environment variable:

```
$ echo $PATH
```

Remove any Oracle environment variables.

11. If you are not installing the software on the local system, then enter a command similar to the following to direct X applications to display on the local system:

- Bourne, Bash, or Korn shell:

```
$ export DISPLAY=local_host:0.0
```
- C shell:

```
% setenv DISPLAY local_host:0.0
```

In this example, `local_host` is the host name or IP address of the system (your workstation, or another client) on which you want to display the installer.

12. If the `/tmp` directory has less than 1 GB of free space, then identify a file system with at least 1 GB of free space and set the `TMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:

Note:

You cannot use a shared file system as the location of the temporary file directory (typically `/tmp`) for Oracle RAC installations. If you place `/tmp` on a shared file system, then the installation fails.

- a. Use the `df -h` command to identify a suitable file system with sufficient free space.
- b. If necessary, enter commands similar to the following to create a temporary directory on the file system that you identified, and set the appropriate permissions on the directory:

```
$ sudo - s  
# mkdir /mount_point/tmp  
# chmod 775 /mount_point/tmp  
# exit
```

- c. Enter commands similar to the following to set the `TMP` and `TMPDIR` environment variables:

Bourne, Bash, or Korn shell:

```
$ TMP=/mount_point/tmp
$ TMPDIR=/mount_point/tmp
$ export TMP TMPDIR
```

C shell:

```
% setenv TMP /mount_point/tmp
% setenv TMPDIR /mount_point/tmp
```

13. To verify that the environment has been set correctly, enter the following commands:

```
$ umask
$ env | more
```

Verify that the `umask` command displays a value of 22, 022, or 0022 and that the environment variables you set in this section have the correct values.

6.5.3 Checking Resource Limits for Oracle Software Installation Users

For each installation software owner user account, check the resource limits for installation.

On Oracle Linux systems, Oracle recommends that you install Oracle Preinstallation RPMs to meet preinstallation requirements like configuring your operating system to set the resource limits in the `limits.conf` file. Oracle Preinstallation RPM only configures the `limits.conf` file for the `oracle` user. If you are implementing Oracle Grid Infrastructure job role separation, then copy the values from the `oracle` user to the `grid` user in the `limits.conf` file.

Use the following ranges as guidelines for resource allocation to Oracle installation owners:

Table 6-1 Installation Owner Resource Limit Recommended Ranges

Resource Shell Limit	Resource	Soft Limit	Hard Limit
Open file descriptors	<code>nofile</code>	at least 1024	at least 65536
Number of processes available to a single user	<code>nproc</code>	at least 2047	at least 16384
Size of the stack segment of the process	<code>stack</code>	at least 10240 KB	at least 10240 KB, and at most 32768 KB
Maximum locked memory limit	<code>memlock</code>	at least 90 percent of the current RAM when HugePages memory is enabled and at least 3145728 KB (3 GB) when HugePages memory is disabled	at least 90 percent of the current RAM when HugePages memory is enabled and at least 3145728 KB (3 GB) when HugePages memory is disabled

To check resource limits:

1. Log in as an installation owner.

2. Check the soft and hard limits for the file descriptor setting. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Sn
1024
$ ulimit -Hn
65536
```

3. Check the soft and hard limits for the number of processes available to a user. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Su
2047
$ ulimit -Hu
16384
```

4. Check the soft limit for the stack setting. Ensure that the result is in the recommended range. For example:

```
$ ulimit -Ss
10240
$ ulimit -Hs
32768
```

5. Repeat this procedure for each Oracle software installation owner.

If necessary, update the resource limits in the `/etc/security/limits.conf` configuration file for the installation owner. However, the configuration file may be distribution specific. Contact your system administrator for distribution specific configuration file information.

Note:

If you make changes to an Oracle installation user account and that user account is logged in, then changes to the `limits.conf` file do not take effect until you log these users out and log them back in. You must do this before you use these accounts for installation.

6.5.4 Setting Remote Display and X11 Forwarding Configuration

If you are on a remote terminal, and the local system has only one visual (which is typical), then use the following syntax to set your user account `DISPLAY` environment variable:

Remote Display

Bourne, Korn, and Bash shells

```
$ export DISPLAY=hostname:0
```

C shell

```
$ setenv DISPLAY hostname:0
```

For example, if you are using the Bash shell and if your host name is `local_host`, then enter the following command:

```
$ export DISPLAY=node1:0
```


X11 Forwarding

To ensure that X11 forwarding does not cause the installation to fail, use the following procedure to create a user-level SSH client configuration file for Oracle installation owner user accounts:

1. Using any text editor, edit or create the software installation owner's `~/.ssh/config` file.
2. Ensure that the `ForwardX11` attribute in the `~/.ssh/config` file is set to `no`. For example:

```
Host *
    ForwardX11 no
```

3. Ensure that the permissions on `~/.ssh` are secured to the Oracle installation owner user account. For example:

```
$ ls -al .ssh
total 28
drwx----- 2 grid oinstall 4096 Jun 21 2015
drwx----- 19 grid oinstall 4096 Jun 21 2015
-rw-r--r-- 1 grid oinstall 1202 Jun 21 2015 authorized_keys
-rwx----- 1 grid oinstall 668 Jun 21 2015 id_dsa
-rwx----- 1 grid oinstall 601 Jun 21 2015 id_dsa.pub
-rwx----- 1 grid oinstall 1610 Jun 21 2015 known_hosts
```

6.5.5 Preventing Installation Errors Caused by Terminal Output Commands

During an Oracle Grid Infrastructure installation, OUI uses SSH to run commands and copy files to the other nodes. During the installation, hidden files on the system (for example, `.bashrc` or `.cshrc`) can cause `makefile` and other installation errors if they contain terminal output commands.

To avoid this problem, you must modify hidden files in each Oracle installation owner user home directory to suppress all output on `STDOUT` or `STDERR` (for example, `stty`, `xtitle`, and other such commands) as in the following examples:

Bourne, Bash, or Korn shell:

```
if [ -t 0 ]; then
    stty intr ^C
fi
```

C shell:

```
test -t 0
if ($status == 0) then
    stty intr ^C
endif
```

Note:

If the remote shell can load hidden files that contain `stty` commands, then OUI indicates an error and stops the installation.

6.6 Enabling Intelligent Platform Management Interface (IPMI)

Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to computer hardware and firmware that system administrators can use to monitor system health and manage the system.

Oracle Clusterware can integrate IPMI to provide failure isolation support and to ensure cluster integrity. You can configure node-termination with IPMI during installation by selecting IPMI from the Failure Isolation Support screen. You can also configure IPMI after installation with `crsctl` commands.

[Requirements for Enabling IPMI](#) (page 6-28)

You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

[Configuring the IPMI Management Network](#) (page 6-29)

You can configure the BMC for DHCP, or for static IP addresses. Oracle recommends that you configure the BMC for dynamic IP address assignment using DHCP. To use this option, you must have a DHCP server configured to assign the BMC IP addresses.

[Configuring the Open IPMI Driver](#) (page 6-29)

On Linux systems, the OpenIPMI driver is the supported driver for Oracle Clusterware deployments with IPMI. You can install and configure the driver dynamically by manually loading the required modules.

[Configuring the BMC](#) (page 6-30)

Configure BMC on each node for remote control using LAN for IPMI-based node fencing to function properly.

[Example of BMC Configuration Using IPMITool](#) (page 6-31)

The following is an example of configuring BMC using `ipmitool` (version 1.8.6).

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to configure IPMI after installation.

6.6.1 Requirements for Enabling IPMI

You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

- Each cluster member node requires a Baseboard Management Controller (BMC) running firmware compatible with IPMI version 1.5 or greater, which supports IPMI over LANs, and configured for remote control using LAN.
- The cluster requires a management network for IPMI. This can be a shared network, but Oracle recommends that you configure a dedicated network.
- Each cluster member node's port used by BMC must be connected to the IPMI management network.
- Each cluster member must be connected to the management network.

- Some server platforms put their network interfaces into a power saving mode when they are powered off. In this case, they may operate only at a lower link speed (for example, 100 MB, instead of 1 GB). For these platforms, the network switch port to which the BMC is connected must be able to auto-negotiate down to the lower speed, or IPMI will not function properly.

Note:

IPMI operates on the physical hardware platform through the network interface of the baseboard management controller (BMC). Depending on your system configuration, an IPMI-initiated restart of a server can affect all virtual environments hosted on the server. Contact your hardware and OS vendor for more information.

6.6.2 Configuring the IPMI Management Network

You can configure the BMC for DHCP, or for static IP addresses. Oracle recommends that you configure the BMC for dynamic IP address assignment using DHCP. To use this option, you must have a DHCP server configured to assign the BMC IP addresses.

For Oracle Clusterware to communicate with the BMC, the IPMI driver must be installed permanently on each node, so that it is available on system restarts. The IPMI driver is available on the Asianux Linux, Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server distributions supported with this release.

Configuring the BMC with dynamic addresses (DHCP) is not supported on HP-UX and AIX platforms.

Note:

If you configure IPMI, and you use Grid Naming Service (GNS) you still must configure separate addresses for the IPMI interfaces. As the IPMI adapter is not seen directly by the host, the IPMI adapter is not visible to GNS as an address on the host.

6.6.3 Configuring the Open IPMI Driver

On Linux systems, the OpenIPMI driver is the supported driver for Oracle Clusterware deployments with IPMI. You can install and configure the driver dynamically by manually loading the required modules.

Contact your Linux distribution vendor for information about how to configure IPMI for your distribution. Configure the Open IPMI driver manually on Oracle Linux as follows:

1. Log in as `root`.

2. Run the following commands:

```
# /sbin/modprobe ipmi_msghandler
# /sbin/modprobe ipmi_si
# /sbin/modprobe ipmi_devintf
```

3. (Optional) Run the command `/sbin/lsmmod |grep ipmi` to confirm that the IPMI modules are loaded. For example:

```
# /sbin/lsmmod | grep ipmi
ipmi_devintf 12617 0
ipmi_si 33377 0
ipmi_msghandler 33701 2
ipmi_devintf,ipmi_si
```

You can install the modules whether or not a BMC is present.

4. Open the `/etc/rc.local` file using a text editor, navigate to the end of the file, and enter lines similar to the following, to run the `modprobe` commands in step 2 automatically on system restart:

```
# START IPMI ON SYSTEM RESTART
/sbin/modprobe ipmi_msghandler
/sbin/modprobe ipmi_si
/sbin/modprobe ipmi_devintf
```

On SUSE Linux Enterprise Server systems, add the `modprobe` commands above to `/etc/init.d/boot.local`.

5. Check to ensure that the Linux system is recognizing the IPMI device, using the following command:

```
ls -l /dev/ipmi0
```

If the IPMI device is dynamically loaded, then the output must be similar to the following:

```
# ls -l /dev/ipmi0 crw----- 1 root root 253, 0 Sep 23 06:29 /dev/ipmi0
```

If you do see the device file output, then the IPMI driver is configured, and you can ignore the following step.

If you do not see the device file output, then the `udev` daemon is not set up to create device files automatically. Proceed to the next step.

6. Determine the device major number for the IPMI device using the command `grep ipmi /proc/devices`. For example:

```
# grep ipmi /proc/devices 253 ipmidev
```

The device major number is 253.

7. Run the `mknod` command to create a directory entry and `i-node` for the IPMI device, using the device major number. For example:

```
# mknod /dev/ipmi0 c 253 0x0
```

The permissions on `/dev/ipmi0` in the example allow the device to be accessible only by root. The device should only be accessed by root, to prevent a system vulnerability.

6.6.4 Configuring the BMC

Configure BMC on each node for remote control using LAN for IPMI-based node fencing to function properly.

You can configure BMC from the BIOS prompt, using a distribution-specific management utility, or you can configure BMC using publicly available utilities, such as the following:

IPMITool, which is available for Linux:

<http://ipmitool.sourceforge.net>

IPMIutil, which is available for Linux and Windows:

<http://ipmiutil.sourceforge.net>

Refer to the documentation for the configuration tool you select for details about using the tool to configure the BMC.

When you configure the BMC on each node, you must complete the following:

1. Enable IPMI over LAN, so that the BMC can be controlled over the management network.
2. Enable dynamic IP addressing using DHCP or GNS, or configure a static IP address for the BMC.
3. Establish an administrator user account and password for the BMC.
4. Configure the BMC for VLAN tags, if you will use the BMC on a tagged VLAN.

The configuration tool you use does not matter, but these conditions must be met for the BMC to function properly.

6.6.5 Example of BMC Configuration Using IPMITool

The following is an example of configuring BMC using `ipmitool` (version 1.8.6).

1. Log in as `root`.
2. Verify that `ipmitool` can communicate with the BMC using the IPMI driver by using the command `bmc info`, and looking for a device ID in the output. For example:

```
# ipmitool bmc info
Device ID           : 32
.
.
.
```

If `ipmitool` is not communicating with the BMC, then review the section “Configuring the BMC” and ensure that the IPMI driver is running.

3. Enable IPMI over LAN using the following procedure
 - a. Determine the channel number for the channel used for IPMI over LAN. Beginning with channel 1, run the following command until you find the channel that displays LAN attributes (for example, the IP address):

```
# ipmitool lan print 1

. . .
IP Address Source   : 0x01
IP Address          : 140.87.155.89
. . .
```

- b. Turn on LAN access for the channel found. For example, where the channel is 1:

```
# ipmitool -I bmc lan set 1 access on
```

4. Configure IP address settings for IPMI using the static IP addressing procedure:

- **Using static IP Addressing**

If the BMC shares a network connection with ILOM, then the IP address must be on the same subnet. You must set not only the IP address, but also the proper values for netmask, and the default gateway. For example, assuming the channel is 1:

```
# ipmitool -I bmc lan set 1 ipaddr 192.168.0.55
# ipmitool -I bmc lan set 1 netmask 255.255.255.0
# ipmitool -I bmc lan set 1 defgw ipaddr 192.168.0.1
```

Note that the specified address (192.168.0.55) is associated only with the BMC, and does not respond to normal pings.

5. Establish an administration account with a username and password, using the following procedure (assuming the channel is 1):

a. Set BMC to require password authentication for ADMIN access over LAN. For example:

```
# ipmitool -I bmc lan set 1 auth ADMIN MD5,PASSWORD
```

b. List the account slots on the BMC, and identify an unused slot less than the maximum ID and not listed, for example, ID 4 in the following example. Note that some slots may be reserved and not available for reuse on some hardware.

```
# ipmitool user summary 1
Maximum IDs      : 20
Enabled User Count : 3
Fixed Name Count  : 2
# ipmitool user list 1
ID Name      Enabled Callin Link Auth IPMI Msg Channel Priv Lim
1          true   false  false   true   USER
2  root      true   false  false   true   ADMINISTRATOR
3  sysoper   true   true   false   true   OPERATOR
12 default   true   true   false   true   NO ACCESS
13          true   false  true    false  CALLBACK
```

In the example above, there are 20 possible slots, and the first unused slot is number 4.

c. Assign the desired administrator user name and password and enable messaging for the identified slot. (Note that for IPMI v1.5 the user name and password can be at most 16 characters). Also, set the privilege level for that slot when accessed over LAN (channel 1) to ADMIN (level 4). For example, where *username* is the administrative user name, and *password* is the password:

```
# ipmitool user set name 4 username
# ipmitool user set password 4 password
# ipmitool user enable 4
# ipmitool channel setaccess 1 4 privilege=4
# ipmitool channel setaccess 1 4 link=on
# ipmitool channel setaccess 1 4 ipmi=on
```

d. Verify the setup using the command `lan print 1`. The output should appear similar to the following. Note that the items in **bold** text are the

settings made in the preceding configuration steps, and comments or alternative options are indicated within brackets []:

```
# ipmitool lan print 1
Set in Progress      : Set Complete
Auth Type Support    : NONE MD2 MD5 PASSWORD
Auth Type Enable     : Callback : MD2 MD5
                    : User       : MD2 MD5
                    : Operator  : MD2 MD5
                    : Admin     : MD5 PASSWORD
                    : OEM       : MD2 MD5
IP Address Source   : DHCP Address [or Static Address]
IP Address         : 192.168.0.55
Subnet Mask       : 255.255.255.0
MAC Address          : 00:14:22:23:fa:f9
SNMP Community String : public
IP Header            : TTL=0x40 Flags=0x40 Precedence=...
Default Gateway IP : 192.168.0.1
Default Gateway MAC  : 00:00:00:00:00:00
.
.
.
# ipmitool channel getaccess 1 4
Maximum User IDs     : 10
Enabled User IDs     : 2

User ID              : 4
User Name         : username [This is the administration user]
Fixed Name           : No
Access Available     : call-in / callback
Link Authentication : enabled
IPMI Messaging     : enabled
Privilege Level    : ADMINISTRATOR
```

6. Verify that the BMC is accessible and controllable from a remote node in your cluster using the `bmc info` command. For example, if `node2-ipmi` is the network host name assigned the IP address of node2's BMC, then to verify the BMC on node `node2` from node1, with the administrator account username, enter the following command on node1:

```
$ ipmitool -H node2-ipmi -U username lan print 1
```

You are prompted for a password. Provide the IPMI password.

If the BMC is correctly configured, then you should see information about the BMC on the remote node. If you see an error message, such as `Error: Unable to establish LAN session`, then you must check the BMC configuration on the remote node.

7. Repeat this process for each cluster member node. If the IPMI administrator account credentials on each cluster member node are not identical, then IPMI will fail during configuration.

6.7 Determining Root Script Execution Plan

During Oracle Grid Infrastructure installation, the installer requires you to run scripts with superuser (or `root`) privileges to complete a number of system configuration tasks.

You can continue to run scripts manually as `root`, or you can delegate to the installer the privilege to run configuration steps as `root`, using one of the following options:

- Use the `root` password: Provide the password to the installer as you are providing other configuration information. The password is used during installation, and not stored. The root user password must be identical on each cluster member node.

To enable root command delegation, provide the `root` password to the installer when prompted.

- Use Sudo: Sudo is a UNIX and Linux utility that allows members of the sudoers list privileges to run individual commands as `root`.

To enable Sudo, have a system administrator with the appropriate privileges configure a user that is a member of the sudoers list, and provide the user name and password when prompted during installation.

Supported Storage Options for Oracle Database and Oracle Grid Infrastructure

Review supported storage options as part of your installation planning process.

[Supported Storage Options for Oracle Grid Infrastructure](#) (page 7-2)

The following table shows the storage options supported for Oracle Grid Infrastructure binaries and files:

[Oracle ACFS and Oracle ADVM](#) (page 7-4)

This section contains information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM).

[Storage Considerations for Oracle Grid Infrastructure and Oracle RAC](#) (page 7-6)

For all installations, you must choose the storage option to use for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters (Oracle RAC) databases.

[Guidelines for Using Oracle ASM Disk Groups for Storage](#) (page 7-7)

Plan how you want to configure Oracle ASM disk groups for deployment.

[Guidelines for Configuring Oracle ASM Disk Groups on NFS](#) (page 7-8)

During installation of Oracle Grid Infrastructure, you can choose to configure Oracle ASM disk groups on NFS file systems on a supported Network Attached Storage (NAS) device.

[Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC](#) (page 7-10)

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers.

[Using a Cluster File System for Oracle Clusterware Files](#) (page 7-10)

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), you must use Oracle Automatic Storage Management (Oracle ASM) to store voting files and OCR files.

[About NFS Storage for Data Files](#) (page 7-10)

Review this section for NFS storage configuration guidelines.

[About Direct NFS Client Mounts to NFS Storage Devices](#) (page 7-11)

Direct NFS Client integrates the NFS client functionality directly in the Oracle software to optimize the I/O path between Oracle and the NFS server. This integration can provide significant performance improvements.

7.1 Supported Storage Options for Oracle Grid Infrastructure

The following table shows the storage options supported for Oracle Grid Infrastructure binaries and files:

Table 7-1 Supported Storage Options for Oracle Grid Infrastructure

Storage Option	OCR and Voting Files	Oracle Clusterware Binaries	Oracle RAC Database Binaries	Oracle RAC Database Data Files	Oracle RAC Database Recovery Files
Oracle Automatic Storage Management (Oracle ASM) Note: Loopback devices are not supported for use with Oracle ASM	Yes	No	No	Yes	Yes
Oracle Automatic Storage Management Cluster File System (Oracle ACFS)	No	No	Yes for Oracle Database 11g Release 2 (11.2) and for Hub Nodes for Oracle Database 12c Release 1 (12.1) and later. No for running Oracle Database on Leaf Nodes.	Yes for Oracle Database 12c Release 1 (12.1) and later	Yes for Oracle Database 12c Release 1 (12.1) and later
Local file system	No	Yes	Yes	No	No
OCFS2	No	No	Yes	Yes	Yes

Table 7-1 (Cont.) Supported Storage Options for Oracle Grid Infrastructure

Storage Option	OCR and Voting Files	Oracle Clusterware Binaries	Oracle RAC Database Binaries	Oracle RAC Database Data Files	Oracle RAC Database Recovery Files
Network file system (NFS) on a certified network-attached storage (NAS) filer Note: Direct NFS Client does not support Oracle Clusterware files	No	Yes	Yes	Yes	Yes
Direct-attached storage (DAS)	No	No	Yes	Yes	Yes
Shared disk partitions (block devices or raw devices)	No	No	No	No	No

Guidelines for Storage Options

Use the following guidelines when choosing storage options:

- You can choose any combination of the supported storage options for each file type provided that you satisfy all requirements listed for the chosen storage options.
- You can only use Oracle ASM to store Oracle Clusterware files.
- Direct use of raw or block devices is not supported. You can only use raw or block devices under Oracle ASM.

See:

Oracle Database Upgrade Guide for information about how to prepare for upgrading an existing database

Note: For information about OCFS2, see the following website:

<http://oss.oracle.com/projects/ocfs2/>

For OCFS2 certification status, and for other cluster file system support, see the Certify page on My Oracle Support.

7.2 Oracle ACFS and Oracle ADVM

This section contains information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM).

Oracle ACFS extends Oracle ASM technology to support of all of your application data in both single instance and cluster configurations. Oracle ADVM provides volume management services and a standard disk device driver interface to clients. Oracle Automatic Storage Management Cluster File System communicates with Oracle ASM through the Oracle Automatic Storage Management Dynamic Volume Manager interface.

[Oracle ACFS and Oracle ADVM Support on Linux x86-64](#) (page 7-4)

Oracle ACFS and Oracle ADVM are supported on Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server.

[Restrictions and Guidelines for Oracle ACFS](#) (page 7-5)

Review these topics as part of your storage plan for using Oracle ACFS for single instance and cluster configurations.

Related Topics:

Oracle Automatic Storage Management Administrator's Guide

7.2.1 Oracle ACFS and Oracle ADVM Support on Linux x86-64

Oracle ACFS and Oracle ADVM are supported on Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server.

Table 7-2 Platforms That Support Oracle ACFS and Oracle ADVM

Platform / Operating System	Kernel
Oracle Linux 7	<ul style="list-style-type: none"> Oracle Linux 7 with Red Hat Compatible Kernel Oracle Linux 7 with the Unbreakable Enterprise Kernel: 3.8.13-35 and later UEK 3.8.13 kernels Oracle Linux 7 with the Unbreakable Enterprise Kernel: 4.1.12-32 and later 4.1.12 kernels
Oracle Linux 6	<ul style="list-style-type: none"> Oracle Linux 6 with Red Hat Compatible Kernel Oracle Linux 6 with the Unbreakable Enterprise Kernel Release 2: 2.6.39-100 and later UEK 2.6.39 kernels Oracle Linux 6 with the Unbreakable Enterprise Kernel Release 3: 3.8.13 and later UEK 3.8.13 kernels Oracle Linux 6 with Unbreakable Enterprise Kernel: 4.1.12-32 and later 4.1.12 kernels
Red Hat Enterprise Linux 7	All Red Hat Enterprise Linux compatible kernels
Red Hat Enterprise Linux 6	All Red Hat Enterprise Linux compatible kernels

Table 7-2 (Cont.) Platforms That Support Oracle ACFS and Oracle ADVM

Platform / Operating System	Kernel
SUSE Linux Enterprise Server 12	All SUSE Linux Enterprise Server 12 kernels
Linux Containers	Not supported

Note:

If you use Security Enhanced Linux (SELinux) in enforcing mode with Oracle ACFS, then ensure that you mount the Oracle ACFS file systems with an SELinux default context. Refer to your Linux vendor documentation for information about the context mount option.

See Also:

- My Oracle Support Note 1369107.1 for more information and a complete list of platforms and releases that support Oracle ACFS and Oracle ADVM:
<https://support.oracle.com/rs?type=doc&id=1369107.1>
- Patch Set Updates for Oracle Products (My Oracle Support Note 854428.1) for current release and support information:
<https://support.oracle.com/rs?type=doc&id=854428.1>

7.2.2 Restrictions and Guidelines for Oracle ACFS

Review these topics as part of your storage plan for using Oracle ACFS for single instance and cluster configurations.

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS) provides a general purpose file system.
- You can only use Oracle ACFS when Oracle ASM is configured.
- Note the following general guidelines and restrictions for placing Oracle Database and Oracle Grid Infrastructure files on Oracle ACFS:
 - You can place Oracle Database binaries, data files, and administrative files (for example, trace files) on Oracle ACFS.
 - Oracle ACFS does not support replication or encryption with Oracle Database data files, tablespace files, control files, and redo logs.
 - You can place Oracle Database homes on Oracle ACFS only if the database release is Oracle Database 11g Release 2, or more recent releases. You cannot install earlier releases of Oracle Database on Oracle ACFS.
 - For installations on Oracle Clusterware, you cannot place Oracle Clusterware files on Oracle ACFS.

- For policy-managed Oracle Flex Cluster databases, Oracle ACFS can run on Hub Nodes, but cannot run on Leaf Nodes. For this reason, Oracle RAC binaries cannot be placed on Oracle ACFS located on Leaf Nodes.
- Oracle Restart does not support root-based Oracle Clusterware resources. For this reason, the following restrictions apply if you run Oracle ACFS on an Oracle Restart Configuration:
 - Oracle Restart does not support Oracle ACFS resources on all platforms.
 - Starting with Oracle Database 12c, Oracle Restart configurations do not support the Oracle ACFS registry.
 - You must manually load Oracle ACFS drivers after a system restart.
 - You must manually mount an Oracle ACFS file system, and unmount it after the Oracle ASM instance has finished running.
 - Creating Oracle data files on an Oracle ACFS file system is not supported in Oracle Restart configurations. Creating Oracle data files on an Oracle ACFS file system is supported on Oracle Grid Infrastructure for a cluster configurations.
- Oracle ACFS and Oracle ADVN are not supported on IBM AIX Workload Partitions (WPARs).

Related Topics:

Oracle Automatic Storage Management Administrator's Guide

7.3 Storage Considerations for Oracle Grid Infrastructure and Oracle RAC

For all installations, you must choose the storage option to use for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters (Oracle RAC) databases.

Storage Considerations for Oracle Clusterware

Oracle Clusterware voting files are used to monitor cluster node status, and Oracle Cluster Registry (OCR) files contain configuration information about the cluster. You must store Oracle Cluster Registry (OCR) and voting files in Oracle ASM disk groups. You can also store a backup of the OCR file in a disk group. Storage must be shared; any node that does not have access to an absolute majority of voting files (more than half) is restarted.

If you use Oracle ASM disk groups created on Network File System (NFS) for storage, then ensure that you follow the recommendations for mounting NFS described in the topic *Guidelines for Configuring Oracle ASM Disk Groups on NFS*.

Storage Considerations for Oracle RAC

Oracle ASM is a supported storage option for database and recovery files. For all installations, Oracle recommends that you create at least two separate Oracle ASM disk groups: One for Oracle Database data files, and one for recovery files. Oracle recommends that you place the Oracle Database disk group and the recovery files disk group in separate failure groups.

- If you do not use Oracle ASM for database files, then Oracle recommends that you place the data files and the Fast Recovery Area in shared storage located outside

of the Oracle home, in separate locations, so that a hardware failure does not affect availability.

- You can choose any combination of the supported storage options for each file type provided that you satisfy all requirements listed for the chosen storage options.
- If you plan to install an Oracle RAC home on a shared OCFS2 location, then you must upgrade OCFS2 to at least version 1.4.1, which supports shared writable mmmaps.
- To use Oracle ASM with Oracle RAC, and if you are configuring a new Oracle ASM instance, then your system must meet the following conditions:
 - All nodes on the cluster have Oracle Clusterware and Oracle ASM 12c Release 2 (12.2) installed as part of an Oracle Grid Infrastructure for a cluster installation.
 - Any existing Oracle ASM instance on any node in the cluster is shut down.
 - To provide voting file redundancy, one Oracle ASM disk group is sufficient. The Oracle ASM disk group provides three or five copies.

You can use NFS, with or without Direct NFS, to store Oracle Database data files. You cannot use NFS as storage for Oracle Clusterware files.

Related Topics:

[Guidelines for Configuring Oracle ASM Disk Groups on NFS](#) (page 7-8)

During installation of Oracle Grid Infrastructure, you can choose to configure Oracle ASM disk groups on NFS file systems on a supported Network Attached Storage (NAS) device.

7.4 Guidelines for Using Oracle ASM Disk Groups for Storage

Plan how you want to configure Oracle ASM disk groups for deployment.

During Oracle Grid Infrastructure installation, you can create one or two Oracle ASM disk groups. After the Oracle Grid Infrastructure installation, you can create additional disk groups using Oracle Automatic Storage Management Configuration Assistant (ASMCA), SQL*Plus, or Automatic Storage Management Command-Line Utility (ASMCMD).

Choose to create a second disk group during Oracle Grid Infrastructure installation. The first disk group stores the Oracle Cluster Registry (OCR), voting files, and the Oracle ASM password file. The second disk group stores the Grid Infrastructure Management Repository (GIMR) data files and Oracle Cluster Registry (OCR) backup files. Oracle strongly recommends that you store the OCR backup files in a different disk group from the disk group where you store OCR files. In addition, having a second disk group for GIMR is advisable for performance, availability, sizing, and manageability of storage.

Note:

- You must specify the Grid Infrastructure Management Repository (GIMR) location at the time of installing Oracle Grid Infrastructure. You cannot migrate the GIMR from one disk group to another later.
 - For Oracle Domain Services Clusters, you must configure two separate Oracle ASM disk groups, one for OCR and voting files and the other for the GIMR.
-
-

If you install Oracle Database or Oracle RAC after you install Oracle Grid Infrastructure, then you can either use the same disk group for database files, OCR, and voting files, or you can use different disk groups. If you create multiple disk groups before installing Oracle RAC or before creating a database, then you can do one of the following:

- Place the data files in the same disk group as the Oracle Clusterware files.
- Use the same Oracle ASM disk group for data files and recovery files.
- Use different disk groups for each file type.

If you create only one disk group for storage, then the OCR and voting files, database files, and recovery files are contained in the one disk group. If you create multiple disk groups for storage, then you can place files in different disk groups.

With Oracle Database 11g Release 2 (11.2) and later releases, Oracle Database Configuration Assistant (DBCA) does not have the functionality to create disk groups for Oracle ASM.

See Also:

Oracle Automatic Storage Management Administrator's Guide for information about creating disk groups

7.5 Guidelines for Configuring Oracle ASM Disk Groups on NFS

During installation of Oracle Grid Infrastructure, you can choose to configure Oracle ASM disk groups on NFS file systems on a supported Network Attached Storage (NAS) device.

You can create Oracle ASM disk groups on block devices or NFS (Network File System) on a supported Network Attached Storage (NAS) device. Understand the following guidelines for configuring Oracle ASM disk groups on NFS:

Note:

All storage products must be supported by both your server and storage vendors.

Guidelines for Deploying Oracle ASM Disk Groups Without Quorum Disks

- To use an NFS file system, it must be on a supported NAS device. Log in to My Oracle Support at the following URL, and click Certifications to find the most current information about supported NAS devices:
<https://support.oracle.com/>
- NFS file systems must be mounted and available over NFS mounts before you start installation. Refer to your vendor documentation to complete NFS configuration and mounting.
- Direct NFS requires hard mounts. Hard mounting NFS filers prevents corruption which could occur if the client connection were to time out. If an NFS filer hangs on an I/O operation to a mirrored file, then the database and Oracle ASM cannot failover to the surviving mirror copy. Therefore, Oracle recommends that you use external redundancy when you deploy Oracle ASM disk groups on NFS storage.
- Oracle ASM Filter Driver and Oracle ACFS and Oracle ADVM are not supported with NFS. These features are incompatible because of the nature of the operating system interface for NFS-based storage.
- The performance of Oracle software and databases stored on Oracle ASM disk groups on NFS depends on the performance of the network connection between the Oracle server and the NAS device. Oracle recommends that you connect the server to the NAS device using a private dedicated network connection, which should be Gigabit Ethernet or better.
- You can configure Oracle ASM on NFS when you deploy an Oracle Standalone Cluster configuration.
- You can specify separate NFS locations for Oracle ASM disk groups for Oracle Clusterware files and the Grid Infrastructure Management Repository (GIMR).
- The user account with which you perform the installation (`oracle` or `grid`) must have write permissions to create the files in the path that you specify.
- When you choose Oracle ASM on NFS, you cannot use Oracle Automatic Storage Management Cluster File System (Oracle ACFS) for storage. This cluster configuration cannot be used as a Rapid Home Provisioning Server.

Guidelines for Deploying Oracle ASM Disk Groups With Quorum Disks

- SAN-attached storage or iSCSI-attached devices are the preferred ways to connect to quorum disks. If your standard deployment requires NFS to be used as storage, then use soft mounts for NFS-based Oracle ASM quorum disks and hard mounts for other Oracle ASM disks.
- You can use Direct NFS (dNFS) for storage of Oracle Database data files. dNFS does not support soft mounts, so you cannot use dNFS for quorum failure groups. Alternatively, use kernel-based NFS with a soft mount for NFS storage residing in a quorum failure group.
- The quorum failure group feature in Oracle ASM enables use of NFS storage in an Oracle ASM disk group without requiring a hard mount for NFS storage in the quorum failure group. This capability is useful for Oracle Extended Clusters where a third site is required for establishing quorum.

Related Topics:

[Creating Files on a NAS Device for Use with Oracle Automatic Storage Management](#) (page 8-12)

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

7.6 Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers.

Using Logical Volume Managers

Oracle Grid Infrastructure and Oracle RAC only support cluster-aware volume managers. Some third-party volume managers are not cluster-aware, and so are not supported. To confirm that a volume manager you want to use is supported, click **Certifications** on My Oracle Support to determine if your volume manager is certified for Oracle RAC. My Oracle Support is available at the following URL:

<https://support.oracle.com>

7.7 Using a Cluster File System for Oracle Clusterware Files

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), you must use Oracle Automatic Storage Management (Oracle ASM) to store voting files and OCR files.

For new installations, use Oracle Automatic Storage Management (Oracle ASM) to store voting files and OCR files. For Linux x86 (32-bit), Linux86-64 (64-bit) and Linux Itanium platforms, Oracle provides a cluster file system, OCFS2. However, Oracle does not recommend using OCFS2 for Oracle Clusterware files.

If your database files are stored on a shared file system, then you can continue to use the same for database files, instead of moving them to Oracle ASM storage.

7.8 About NFS Storage for Data Files

Review this section for NFS storage configuration guidelines.

Network-Attached Storage and NFS Protocol

Network-attached storage (NAS) systems use the network file system (NFS) protocol to access files over a network, which enables client servers to access files over networks as easily as to storage devices attached directly to the servers. You can store data files on supported NFS systems. NFS is a shared file system protocol, so NFS can support both single instance and Oracle Real Application Clusters databases.

Note:

The performance of Oracle software and databases stored on NAS devices depends on the performance of the network connection between the servers and the network-attached storage devices. For better performance, Oracle recommends that you connect servers to NAS devices using private dedicated network connections. NFS network connections should use Gigabit Ethernet or better.

Refer to your vendor documentation to complete NFS configuration and mounting.

Requirements for Using NFS Storage

Before you start installation, NFS file systems must be mounted and available to servers.

7.9 About Direct NFS Client Mounts to NFS Storage Devices

Direct NFS Client integrates the NFS client functionality directly in the Oracle software to optimize the I/O path between Oracle and the NFS server. This integration can provide significant performance improvements.

Direct NFS Client supports NFSv3, NFSv4, NFSv4.1, and pNFS protocols to access the NFS server. Direct NFS Client also simplifies, and in many cases automates, the performance optimization of the NFS client configuration for database workloads.

Starting with Oracle Database 12c Release 2, when you enable Direct NFS, you can also enable the Direct NFS dispatcher. The Direct NFS dispatcher consolidates the number of TCP connections that are created from a database instance to the NFS server. In large database deployments, using Direct NFS dispatcher improves scalability and network performance. Parallel NFS deployments also require a large number of connections. Hence, the Direct NFS dispatcher is recommended with Parallel NFS deployments too.

Direct NFS Client can obtain NFS mount points either from the operating system mount entries, or from the `oranfstab` file.

Direct NFS Client Requirements

- NFS servers must have write size values (`wtmax`) of 32768 or greater to work with Direct NFS Client.
- NFS mount points must be mounted both by the operating system kernel NFS client and Direct NFS Client, even though you configure Direct NFS Client to provide file service.

If Oracle Database cannot connect to an NFS server using Direct NFS Client, then Oracle Database connects to the NFS server using the operating system kernel NFS client. When Oracle Database fails to connect to NAS storage through Direct NFS Client, it logs an informational message about the Direct NFS Client connect error in the Oracle alert and trace files.

- Follow standard guidelines for maintaining integrity of Oracle Database files mounted by both operating system NFS and by Direct NFS Client.

Direct NFS Mount Point Search Order

Direct NFS Client searches for mount entries in the following order:

1. `$ORACLE_HOME/dbs/oranfstab`
2. `/etc/oranfstab`
3. `/etc/mtab`

Direct NFS Client uses the first matching entry as the mount point.

Note:

You can have only one active Direct NFS Client implementation for each instance. Using Direct NFS Client on an instance prevents another Direct NFS Client implementation.

See Also:

- *Oracle Database Reference* for information about setting the `enable_dnfs_dispatcher` parameter in the initialization parameter file to enable Direct NFS dispatcher
 - *Oracle Database Performance Tuning Guide* for performance benefits of enabling Parallel NFS and Direct NFS dispatcher
 - *Oracle Automatic Storage Management Administrator's Guide* for guidelines about managing Oracle Database data files created with Direct NFS Client or kernel NFS
-
-

Configuring Storage for Oracle Grid Infrastructure

Complete these procedures to configure Oracle Automatic Storage Management (Oracle ASM) for Oracle Grid Infrastructure for a cluster.

Oracle Grid Infrastructure for a cluster provides system support for Oracle Database. Oracle ASM is a volume manager and a file system for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations. Oracle Automatic Storage Management also supports a general purpose file system for your application needs, including Oracle Database binaries. Oracle Automatic Storage Management is Oracle's recommended storage management solution. It provides an alternative to conventional volume managers and file systems.

Note: Oracle ASM is the supported storage management solution for Oracle Cluster Registry (OCR) and Oracle Clusterware voting files. The OCR is a file that contains the configuration information and status of the cluster. The installer automatically initializes the OCR during the Oracle Clusterware installation. Database Configuration Assistant uses the OCR for storing the configurations for the cluster databases that it creates.

[Configuring Storage for Oracle Automatic Storage Management](#) (page 8-2)
Identify storage requirements and Oracle ASM disk group options.

[Configuring Storage Device Path Persistence Using Oracle ASMFD](#) (page 8-13)
Oracle ASM Filter Driver (Oracle ASMFD) maintains storage file path persistence and helps to protect files from accidental overwrites.

[Using Disk Groups with Oracle Database Files on Oracle ASM](#) (page 8-14)
Review this information to configure Oracle Automatic Storage Management (Oracle ASM) storage for Oracle Clusterware and Oracle Database Files.

[Configuring File System Storage for Oracle Database](#) (page 8-16)
Complete these procedures to use file system storage for Oracle Database.

[Creating Member Cluster Manifest File for Oracle Member Clusters](#) (page 8-21)
Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Rapid Home Provisioning configuration.

[Configuring Oracle Automatic Storage Management Cluster File System](#) (page 8-22)
Review this information to configure Oracle ACFS for an Oracle RAC Oracle Database home.

[Checking OCFS2 Version Manually](#) (page 8-23)

Use the following command to check your OCFS2 version manually.

8.1 Configuring Storage for Oracle Automatic Storage Management

Identify storage requirements and Oracle ASM disk group options.

[Identifying Storage Requirements for Oracle Automatic Storage Management](#) (page 8-3)

To identify the storage requirements for using Oracle ASM, you must determine the number of devices and the amount of free disk space that you require.

[Oracle Clusterware Storage Space Requirements](#) (page 8-7)

Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files, and installing the starter database, for various Oracle Cluster deployments.

[About the Grid Infrastructure Management Repository](#) (page 8-10)

Every Oracle Standalone Cluster and Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR), or the Management Database (MGMTDB).

[Using an Existing Oracle ASM Disk Group](#) (page 8-10)

Use Oracle Enterprise Manager Cloud Control or the Oracle ASM command line tool (`asmcmd`) to identify existing disk groups, and to determine if sufficient space is available in the disk group.

[About Upgrading Existing Oracle Automatic Storage Management Instances](#) (page 8-11)

Oracle Automatic Storage Management (Oracle ASM) upgrades are carried out during an Oracle Grid Infrastructure upgrade.

[Selecting Disks to use with Oracle ASM Disk Groups](#) (page 8-11)

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

[Specifying the Oracle ASM Disk Discovery String](#) (page 8-12)

When an Oracle ASM instance is initialized, Oracle ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the `ASM_DISKSTRING` initialization parameter.

[Creating Files on a NAS Device for Use with Oracle Automatic Storage Management](#) (page 8-12)

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

Related Topics:

Oracle Automatic Storage Management Administrator's Guide

8.1.1 Identifying Storage Requirements for Oracle Automatic Storage Management

To identify the storage requirements for using Oracle ASM, you must determine the number of devices and the amount of free disk space that you require.

To complete this task, follow these steps:

1. Plan your Oracle ASM disk groups requirement, based on the cluster configuration you want to deploy. Oracle Domain Services Clusters store Oracle Clusterware files and the Grid Infrastructure Management Repository (GIMR) on separate Oracle ASM disk groups and hence require configuration of two separate Oracle ASM disk groups, one for OCR and voting files and the other for the GIMR.
2. Determine whether you want to use Oracle ASM for Oracle Database files, recovery files, and Oracle Database binaries. Oracle Database files include data files, control files, redo log files, the server parameter file, and the password file.

Note:

- You do not have to use the same storage mechanism for Oracle Database files and recovery files. You can use a shared file system for one file type and Oracle ASM for the other.
 - There are two types of Oracle Clusterware files: OCR files and voting files. You must use Oracle ASM to store OCR and voting files.
 - If your database files are stored on a shared file system, then you can continue to use the same for database files, instead of moving them to Oracle ASM storage.
-
-

3. Choose the Oracle ASM redundancy level to use for the Oracle ASM disk group.

Except when using external redundancy, Oracle ASM mirrors all Oracle Clusterware files in separate failure groups within a disk group. A quorum failure group, a special type of failure group, contains mirror copies of voting files when voting files are stored in normal or high redundancy disk groups. The disk groups that contain Oracle Clusterware files (OCR and voting files) have a higher minimum number of failure groups than other disk groups because the voting files are stored in quorum failure groups in the Oracle ASM disk group.

A quorum failure group is a special type of failure group that is used to store the Oracle Clusterware voting files. The quorum failure group is used to ensure that a quorum of the specified failure groups are available. When Oracle ASM mounts a disk group that contains Oracle Clusterware files, the quorum failure group is used to determine if the disk group can be mounted in the event of the loss of one or more failure groups. Disks in the quorum failure group do not contain user data, therefore a quorum failure group is not considered when determining redundancy requirements in respect to storing user data.

The redundancy levels are as follows:

- High redundancy

In a high redundancy disk group, Oracle ASM uses three-way mirroring to increase performance and provide the highest level of reliability. A high redundancy disk group requires a minimum of three disk devices (or three

failure groups). The effective disk space in a high redundancy disk group is one-third the sum of the disk space in all of its devices.

For Oracle Clusterware files, a high redundancy disk group requires a minimum of five disk devices and provides five voting files and one OCR (one primary and two secondary copies). For example, your deployment may consist of three regular failure groups and two quorum failure groups. Note that not all failure groups can be quorum failure groups, even though voting files need all five disks. With high redundancy, the cluster can survive the loss of two failure groups.

While high redundancy disk groups do provide a high level of data protection, you should consider the greater cost of additional storage devices before deciding to select high redundancy disk groups.

- Normal redundancy

In a normal redundancy disk group, to increase performance and reliability, Oracle ASM by default uses two-way mirroring. A normal redundancy disk group requires a minimum of two disk devices (or two failure groups). The effective disk space in a normal redundancy disk group is half the sum of the disk space in all of its devices.

For Oracle Clusterware files, a normal redundancy disk group requires a minimum of three disk devices and provides three voting files and one OCR (one primary and one secondary copy). For example, your deployment may consist of two regular failure groups and one quorum failure group. With normal redundancy, the cluster can survive the loss of one failure group.

If you are not using a storage array providing independent protection against data loss for storage, then Oracle recommends that you select normal redundancy.

- External redundancy

An external redundancy disk group requires a minimum of one disk device. The effective disk space in an external redundancy disk group is the sum of the disk space in all of its devices.

Because Oracle ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use external redundancy with storage devices such as RAID, or other similar devices that provide their own data protection mechanisms.

- Flex redundancy

A flex redundancy disk group is a type of redundancy disk group with features such as flexible file redundancy, mirror splitting, and redundancy change. A flex disk group can consolidate files with different redundancy requirements into a single disk group. It also provides the capability for databases to change the redundancy of its files. A disk group is a collection of file groups, each associated with one database. A quota group defines the maximum storage space or quota limit of a group of databases within a disk group.

In a flex redundancy disk group, Oracle ASM uses three-way mirroring of Oracle ASM metadata to increase performance and provide reliability. For database data, you can choose no mirroring (unprotected), two-way mirroring (mirrored), or three-way mirroring (high). A flex redundancy disk group requires a minimum of three disk devices (or three failure groups).

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about file groups and quota groups for flex disk groups

Note:

You can alter the redundancy level of the disk group after a disk group is created. For example, you can convert a normal or high redundancy disk group to a flex redundancy disk group. Within a flex redundancy disk group, file redundancy can change among three possible values: unprotected, mirrored, or high.

4. Determine the total amount of disk space that you require for Oracle Clusterware files, and for the database files and recovery files.

If an Oracle ASM instance is running on the system, then you can use an existing disk group to meet these storage requirements. If necessary, you can add disks to an existing disk group during the database installation.

See *Oracle Clusterware Storage Space Requirements* to determine the minimum number of disks and the minimum disk space requirements for installing Oracle Clusterware files, and installing the starter database, where you have voting files in a separate disk group.

5. Determine an allocation unit size.

Every Oracle ASM disk is divided into allocation units (AU). An allocation unit is the fundamental unit of allocation within a disk group. You can select the AU Size value from 1, 2, 4, 8, 16, 32 or 64 MB, depending on the specific disk group compatibility level. For flex disk groups, the default value for AU size is set to 4 MB. For external, normal, and high redundancies, the default AU size is 1 MB.

6. For Oracle Clusterware installations, you must also add additional disk space for the Oracle ASM metadata. You can use the following formula to calculate the disk space requirements (in MB) for OCR and voting files, and the Oracle ASM metadata:

$$\text{total} = [2 * \text{ausize} * \text{disks}] + [\text{redundancy} * (\text{ausize} * (\text{all_client_instances} + \text{nodes} + \text{disks} + 32) + (64 * \text{nodes}) + \text{clients} + 543)]$$

redundancy = Number of mirrors: external = 1, normal = 2, high = 3, flex = 3.

ausize = Metadata AU size in megabytes

nodes = Number of nodes in cluster.

clients - Number of database instances for each node.

disks - Number of disks in disk group.

For example, for a four-node Oracle RAC installation, using three disks in a normal redundancy disk group, you require an additional 5293 MB of space: $[2 * 4 * 3] + [2 * (4 * (4 * (4 + 1) + 30) + (64 * 4) + 533)] = 5293$ MB

7. Optionally, identify failure groups for the Oracle ASM disk group devices.

If you intend to use a normal or high redundancy disk group, then you can further protect the database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device is included in its failure group.

However, if two disk devices in a normal redundancy disk group are attached to the same Host Bus Adapter (HBA), then the disk group becomes unavailable if the adapter fails. The HBA in this example is a single point of failure.

For instance, to avoid failures of this type, you can use two HBA fabric paths, each with two disks, and define a failure group for the disks attached to each adapter. This configuration would enable the disk group to tolerate the failure of one HBA fabric path.

Note:

You can define custom failure groups during installation of Oracle Grid Infrastructure. You can also define failure groups after installation using the GUI tool ASMCA, the command line tool `asmcmd`, or SQL commands. If you define custom failure groups, then you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

8. If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group. Use the following guidelines when identifying appropriate disk devices:
 - The disk devices must be owned by the user performing Oracle Grid Infrastructure installation.
 - All the devices in an Oracle ASM disk group must be the same size and have the same performance characteristics.
 - Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.
 - Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. Oracle recommends that if you choose to use a logical volume manager, then use the logical volume manager to represent a single logical unit number (LUN) without striping or mirroring, so that you can minimize the effect on storage performance of the additional storage layer.
9. If you use Oracle ASM disk groups created on Network File System (NFS) for storage, then ensure that you follow recommendations described in *Guidelines for Configuring Oracle ASM Disk Groups on NFS*.

Related Topics:[Storage Checklist for Oracle Grid Infrastructure](#) (page 1-9)

Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.

[Oracle Clusterware Storage Space Requirements](#) (page 8-7)

Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files, and installing the starter database, for various Oracle Cluster deployments.

Related Topics:

[Guidelines for Configuring Oracle ASM Disk Groups on NFS](#) (page 7-8)

During installation of Oracle Grid Infrastructure, you can choose to configure Oracle ASM disk groups on NFS file systems on a supported Network Attached Storage (NAS) device.

8.1.2 Oracle Clusterware Storage Space Requirements

Use this information to determine the minimum number of disks and the minimum disk space requirements based on the redundancy type, for installing Oracle Clusterware files, and installing the starter database, for various Oracle Cluster deployments.

Total Storage Space for Database Files Required by Redundancy Type

The following tables list the space requirements for Oracle RAC Database data files for multitenant and non-CDB deployments.

Table 8-1 Oracle ASM Disk Space Minimum Requirements for Oracle Database

Redundancy Level	Minimum number of disks	Data Files	Recovery Files	Both File Types
External	1	4.5 GB	12.9 GB	17.4 GB
Normal	2	8.6 GB	25.8 GB	34.4 GB
High	3	12.9 GB	38.7 GB	51.6 GB
Flex	3	12.9 GB	38.7 GB	51.6 GB

Table 8-2 Oracle ASM Disk Space Minimum Requirements for Oracle Database (non-CDB)

Redundancy Level	Minimum number of disks	Data Files	Recovery Files	Both File Types
External	1	2.7 GB	7.8 GB	10.5 GB
Normal	2	5.2 GB	15.6 GB	20.8 GB
High	3	7.8 GB	23.4 GB	31.2 GB
Flex	3	7.8 GB	23.4 GB	31.2 GB

Total Oracle Clusterware Storage Space Required by Oracle Cluster Deployment Type

During installation of an Oracle Standalone Cluster, if you create the MGMT disk group for Grid Infrastructure Management Repository (GIMR), then the installer requires that you use a disk group with at least 35 GB of available space.

Based on the cluster configuration you want to install, the Oracle Clusterware space requirements vary for different redundancy levels. The following tables list the space requirements for each cluster configuration.

Table 8-3 Minimum Space Requirements for Oracle Domain Services Cluster with Four or Fewer Oracle Member Clusters

Cluster Configuration	Redundancy Level	Space Required for DATA Disk Group containing Oracle Clusterware Files (OCR and Voting Files)	Space Required for MGMT Disk Group containing the GIMR and Oracle Clusterware Backup Files	MGMT Disk Group Space Required for Additional Services, If Selected	Total Storage
Two nodes, 4 MB Allocation Unit (AU), one Oracle ASM disk	External	1.4 GB	188 GB	RHP: 100 GB PDB for each Oracle Member Cluster beyond four: 35 GB	189.4 GB for an Oracle Domain Services Cluster with four Oracle Member Clusters
Two nodes, 4 MB Allocation Unit (AU), three Oracle ASM disks	Normal	2.5 GB	376 GB	RHP: 200 GB PDB for each Oracle Member Cluster beyond four: 70 GB	378.5 GB for an Oracle Domain Services Cluster with four Oracle Member Clusters
Two nodes, 4 MB Allocation Unit (AU), five Oracle ASM disks	High	3.6 GB	564 GB	RHP: 300 GB PDB for each Oracle Member Cluster beyond four: 105 GB	567.6 GB for an Oracle Domain Services Cluster with four Oracle Member Clusters
Two nodes, 4 MB Allocation Unit (AU), three Oracle ASM disks	Flex	2.5 GB	376 GB	RHP: 200 GB PDB for each Oracle Member Cluster beyond four: 70 GB	378.5 GB for an Oracle Domain Services Cluster with four Oracle Member Clusters

The storage space calculations assume an Oracle Domain Services Cluster configuration with four Oracle Member Clusters and two nodes. The Rapid Home Provisioning (RHP) size is 100 GB.

Table 8-4 Minimum Space Requirements for Oracle Member Cluster

Cluster Configuration	OCR	Voting Files	Additional Storage Requirements	Total Storage
Oracle Member Clusters with remote Oracle ASM configuration	1.4 GB	300 MB	None, since Oracle Member Cluster uses the remote GIMR and other services from the Oracle Domain Services Cluster	1.7 GB
Oracle Member Clusters with local Oracle ASM configuration.	Same as the requirements for Oracle Standalone Cluster for each redundancy level.	Same as the requirements for Oracle Standalone Cluster for each redundancy level.	Same as the requirements for Oracle Standalone Cluster for each redundancy level.	Same as the requirements for Oracle Standalone Cluster for each redundancy level.

Table 8-5 Minimum Space Requirements for Oracle Standalone Cluster

Cluster Configuration	Redundancy Level	Space Required for DATA Disk Group containing Oracle Clusterware Files (OCR and Voting Files)	Space Required for MGMT Disk Group containing the GIMR and Oracle Clusterware Backup Files	Total Storage
Two nodes, 4 MB Allocation Unit (AU), one Oracle ASM disks	External	1.4 GB	At least 37.6 GB for a cluster with 4 nodes or less. Additional 4.7 GB space required for clusters with 5 or more nodes.	39 GB
Two nodes, 4 MB Allocation Unit (AU), three Oracle ASM disks	Normal	2.5 GB	75.5 GB	78 GB
Two nodes, 4 MB Allocation Unit (AU), five Oracle ASM disks	High	3.6 GB	113.4 GB	117 GB
Two nodes, 4 MB Allocation Unit (AU), three Oracle ASM disks	Flex	2.5 GB	75.5 GB	78 GB

8.1.3 About the Grid Infrastructure Management Repository

Every Oracle Standalone Cluster and Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR), or the Management Database (MGMTDB).

The Grid Infrastructure Management Repository (GIMR) is a multitenant database with a pluggable database (PDB) for the GIMR of each cluster. The GIMR stores the following information about the cluster:

- Real time performance data the Cluster Health Monitor collects
- Fault, diagnosis, and metric data the Cluster Health Advisor collects
- Cluster-wide events about all resources that Oracle Clusterware collects
- CPU architecture data for Quality of Service Management (QoS)
- Metadata required for Rapid Home Provisioning

The Oracle Standalone Cluster locally hosts the GIMR on an Oracle ASM disk group; this GIMR is a multitenant database with a single pluggable database (PDB).

The global GIMR runs in an Oracle Domain Services Cluster. Oracle Domain Services Cluster locally hosts the GIMR in a separate Oracle ASM disk group. Client clusters, such as Oracle Member Cluster for Database, use the remote GIMR located on the Domain Services Cluster. For two-node or four-node clusters, hosting the GIMR for a cluster on a remote cluster reduces the overhead of running an extra infrastructure repository on a cluster. The GIMR for an Oracle Domain Services Cluster is a multitenant database with one PDB, and additional PDB for each member cluster that is added.

When you configure an Oracle Domain Services Cluster, the installer prompts to configure a separate Oracle ASM disk group for the GIMR, with the default name as MGMT.

Related Topics:

[About Oracle Standalone Clusters](#) (page 9-3)

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

[About Oracle Cluster Domain and Oracle Domain Services Cluster](#) (page 9-3)

An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.

[About Oracle Member Clusters](#) (page 9-5)

Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications.

8.1.4 Using an Existing Oracle ASM Disk Group

Use Oracle Enterprise Manager Cloud Control or the Oracle ASM command line tool (`asmcmd`) to identify existing disk groups, and to determine if sufficient space is available in the disk group.

1. Connect to the Oracle ASM instance and start the instance if necessary:

```
$ $ORACLE_HOME/bin/asmcmd
ASMCMD> startup
```

2. Enter one of the following commands to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

```
ASMCMD> lsdg
```

or

```
$ORACLE_HOME/bin/asmcmd -p lsdg
```

3. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
4. If necessary, install or identify the additional disk devices required to meet the storage requirements for your installation.

Note: If you are adding devices to an existing disk group, then Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

8.1.5 About Upgrading Existing Oracle Automatic Storage Management Instances

Oracle Automatic Storage Management (Oracle ASM) upgrades are carried out during an Oracle Grid Infrastructure upgrade.

If you are upgrading from Oracle ASM 11g Release 2 (11.2.0.1) or later, then Oracle ASM is always upgraded with Oracle Grid Infrastructure as part of the upgrade, and Oracle Automatic Storage Management Configuration Assistant (Oracle ASMCA) is started by the root scripts during upgrade. Subsequently, you can use Oracle ASMCA (located in `Grid_home/bin`) to configure failure groups, Oracle ASM volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

Oracle ASMCA cannot perform a separate upgrade of Oracle ASM from a prior release to the current release.

Upgrades of Oracle ASM from releases prior to 11g Release 2 (11.2) are not supported.

Related Topics:

Oracle Automatic Storage Management Administrator's Guide

Oracle Database Upgrade Guide

8.1.6 Selecting Disks to use with Oracle ASM Disk Groups

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

Use the following guidelines when identifying appropriate disk devices:

- All of the devices in an Oracle ASM disk group should be the same size and have the same performance characteristics.
- Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.

- Nonshared logical partitions are not supported with Oracle RAC. To use logical partitions for your Oracle RAC database, you must use shared logical volumes created by a logical volume manager such as `fdisk`.
- Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

8.1.7 Specifying the Oracle ASM Disk Discovery String

When an Oracle ASM instance is initialized, Oracle ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the `ASM_DISKSTRING` initialization parameter.

The value for the `ASM_DISKSTRING` initialization parameter is an operating system–dependent value that Oracle ASM uses to limit the set of paths that the discovery process uses to search for disks. The exact syntax of a discovery string depends on the platform, ASMLib libraries, and whether Oracle Exadata disks are used. The path names that an operating system accepts are always usable as discovery strings.

The default value of `ASM_DISKSTRING` might not find all disks in all situations. If your site is using a third-party vendor ASMLib, then the vendor might have discovery string conventions that you must use for `ASM_DISKSTRING`. In addition, if your installation uses multipathing software, then the software might place pseudo-devices in a path that is different from the operating system default.

See Also:

- *Oracle Automatic Storage Management Administrator's Guide* for more information about the initialization parameter `ASM_DISKSTRING`
 - See "Oracle ASM and Multipathing" in *Oracle Automatic Storage Management Administrator's Guide* for information about configuring Oracle ASM to work with multipathing, and consult your multipathing vendor documentation for details.
-
-

8.1.8 Creating Files on a NAS Device for Use with Oracle Automatic Storage Management

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Oracle ASM disk group.

Ensure that you specify the ASM discovery path for Oracle ASM disks.

During installation of Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Universal Installer (OUI) can create files in the NFS mounted directory you specify. The following procedure explains how to manually create files in an NFS mounted directory to use as disk devices in an Oracle ASM disk group.

1. If necessary, create an exported directory for the disk group files on the NAS device.
2. Switch user to `root`.

3. Create a mount point directory on the local system.

For example:

```
# mkdir -p /mnt/oracleasm
```

4. To ensure that the NFS file system is mounted when the system restarts, add an entry for the file system in the mount file `/etc/fstab`.
5. Enter a command similar to the following to mount the NFS on the local system:

```
# mount /mnt/oracleasm
```

6. Choose a name for the disk group to create, and create a directory for the files on the NFS file system, using the disk group name as the directory name.

For example, if you want to set up a disk group for a sales database:

```
# mkdir /mnt/oracleasm/sales1
```

7. Use commands similar to the following to create the required number of zero-padded files in this directory:

```
# dd if=/dev/zero
of=/mnt/oracleasm/sales1/disk1 bs=1024k
count=1000
```

This example creates 1 GB files on the NFS file system. You must create one, two, or three files respectively to create an external, normal, or high redundancy disk group.

Note: Creating multiple zero-padded files on the same NAS device does not guard against NAS failure. Instead, create one file for each NAS device and mirror them using the Oracle ASM technology.

8. Enter commands similar to the following to change the owner, group, and permissions on the directory and files that you created:

```
# chown -R grid:asmadmin /mnt/oracleasm
# chmod -R 660 /mnt/oracleasm
```

In this example, the installation owner is `grid` and the OSASM group is `asmadmin`.

9. During Oracle Database installations, edit the Oracle ASM disk discovery string to specify a regular expression that matches the file names you created.

For example:

```
/mnt/oracleasm/sales1/
```

Related Topics:

[My Oracle Support Note 359515.1](#)

8.2 Configuring Storage Device Path Persistence Using Oracle ASMFD

Oracle ASM Filter Driver (Oracle ASMFD) maintains storage file path persistence and helps to protect files from accidental overwrites.

The following references introduce you to Oracle ASMFD:

[About Oracle ASM with Oracle ASM Filter Driver](#) (page 8-14)

During Oracle Grid Infrastructure installation, you can choose to install and configure Oracle Automatic Storage Management Filter Driver (Oracle ASMFD). Oracle ASMFD helps prevent corruption in Oracle ASM disks and files within the disk group.

8.2.1 About Oracle ASM with Oracle ASM Filter Driver

During Oracle Grid Infrastructure installation, you can choose to install and configure Oracle Automatic Storage Management Filter Driver (Oracle ASMFD). Oracle ASMFD helps prevent corruption in Oracle ASM disks and files within the disk group.

Oracle ASM Filter Driver (Oracle ASMFD) rejects write I/O requests that are not issued by Oracle software. This write filter helps to prevent users with administrative privileges from inadvertently overwriting Oracle ASM disks, thus preventing corruption in Oracle ASM disks and files within the disk group. For disk partitions, the area protected is the area on the disk managed by Oracle ASMFD, assuming the partition table is left untouched by the user.

Oracle ASMFD simplifies the configuration and management of disk devices by eliminating the need to rebind disk devices used with Oracle ASM each time the system is restarted.

If Oracle ASMLIB exists on your Linux system, then deinstall Oracle ASMLIB before installing Oracle Grid Infrastructure, so that you can choose to install and configure Oracle ASMFD during an Oracle Grid Infrastructure installation.

Note: Oracle ASMFD is supported on Linux x86-64 and Oracle Solaris operating systems.

Related Topics:

[Deinstalling Oracle ASMLIB On Oracle Grid Infrastructure](#) (page B-14)

If Oracle ASM library driver (Oracle ASMLIB) is installed but you do not use it for device path persistence, then deinstall Oracle ASMLIB.

Oracle Automatic Storage Management Administrator's Guide

8.3 Using Disk Groups with Oracle Database Files on Oracle ASM

Review this information to configure Oracle Automatic Storage Management (Oracle ASM) storage for Oracle Clusterware and Oracle Database Files.

[Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM](#) (page 8-15)

Identify existing disk groups and determine the free disk space that they contain. Optionally, identify failure groups for the Oracle ASM disk group devices.

[Creating Disk Groups for Oracle Database Data Files](#) (page 8-15)

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

[Creating Directories for Oracle Database Files](#) (page 8-15)

Perform this procedure to place the Oracle Database or recovery files on a separate file system from the Oracle base directory.

8.3.1 Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM

Identify existing disk groups and determine the free disk space that they contain. Optionally, identify failure groups for the Oracle ASM disk group devices.

If you intend to use a normal or high redundancy disk group, then you can further protect your database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device comprises its own failure group. However, if two disk devices in a normal redundancy disk group are attached to the same SCSI controller, then the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To protect against failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, then you must specify a minimum of two failure groups for normal redundancy and three failure groups for high redundancy.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for information about Oracle ASM disk discovery.

8.3.2 Creating Disk Groups for Oracle Database Data Files

If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group.

Use the following guidelines when identifying appropriate disk devices:

- All of the devices in an Oracle ASM disk group should be the same size and have the same performance characteristics.
- Do not specify multiple partitions on a single physical disk as a disk group device. Oracle ASM expects each disk group device to be on a separate physical disk.
- Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

8.3.3 Creating Directories for Oracle Database Files

Perform this procedure to place the Oracle Database or recovery files on a separate file system from the Oracle base directory.

1. Use the following command to determine the free disk space on each mounted file system:

```
# df -h
```

- Identify the file systems to use, from the display:

Option	Description
Database Files	Select one of the following: <ul style="list-style-type: none"> A single file system with at least 1.5 GB of free disk space Two or more file systems with at least 3.5 GB of free disk space in total
Recovery Files	Choose a file system with at least 2 GB of free disk space

If you are using the same file system for multiple file types, then add the disk space requirements for each type to determine the total disk space requirement.

- Note the names of the mount point directories for the file systems that you identified.
- If the user performing installation has permissions to create directories on the disks where you plan to install Oracle Database, then DBCA creates the Oracle Database file directory, and the Recovery file directory. If the user performing installation does not have write access, then you must create these directories manually.

For example, given the user `oracle` and Oracle Inventory Group `oinstall`, and using the paths `/u03/oradata/wrk_area` for Oracle Database files, and `/u01/oradata/rcv_area` for the recovery area, these commands create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:

- Database file directory:

```
# mkdir /u01/oradata/
# chown oracle:oinstall /u01/oradata/
# chmod 775 /mount_point/oradata
```

The default location for the database file directory is `$ORACLE_BASE/oradata`.

- Recovery file directory (fast recovery area):

```
# mkdir /u01/oradata/rcv_area
# chown oracle:oinstall /u01/app/oracle/fast_recovery_area
# chmod 775 /u01/oradata/rcv_area
```

The default fast recovery area is `$ORACLE_BASE/fast_recovery_area`.

Oracle recommends that you keep the fast recovery area on a separate physical disk than that of the database file directory. This method enables you to use the fast recovery area to retrieve data if the disk containing `oradata` is unusable for any reason.

8.4 Configuring File System Storage for Oracle Database

Complete these procedures to use file system storage for Oracle Database.

If you plan to place storage on Network File System (NFS) protocol devices, then Oracle recommends that you use Oracle Direct NFS (dNFS) to take advantage of performance optimizations built into the Oracle Direct NFS client.

For optimal database organization and performance, Oracle recommends that you install data files and the Oracle Database software in different disks.

[Configuring NFS Buffer Size Parameters for Oracle Database](#) (page 8-17)

Set the values for the NFS buffer size parameters `rsize` and `wsize` to 32768.

[Checking TCP Network Protocol Buffer for Direct NFS Client](#) (page 8-17)

Check your TCP network buffer size to ensure that it is adequate for the speed of your servers.

[Creating an `orantstab` File for Direct NFS Client](#) (page 8-18)

Direct NFS uses a configuration file, `orantstab`, to determine the available mount points.

[Enabling and Disabling Direct NFS Client Control of NFS](#) (page 8-20)

Use these commands to enable or disable Direct NFS Client Oracle Disk Manager Control of NFS.

[Enabling Hybrid Columnar Compression on Direct NFS Client](#) (page 8-21)

Perform these steps to enable Hybrid Columnar Compression (HCC) on Direct NFS Client:

8.4.1 Configuring NFS Buffer Size Parameters for Oracle Database

Set the values for the NFS buffer size parameters `rsize` and `wsize` to 32768.

For example, to use `rsize` and `wsize` buffer settings with the value 32768 for an Oracle Database data files mount point, set mount point parameters to values similar to the following:

```
nfs_server:/vol/DATA/oradata /home/oracle/netapp nfs\
rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,actimeo=0,vers=3,timeo=600
```

Direct NFS Client issues writes at `wtmax` granularity to the NFS server.

Related Topics:

[My Oracle Support note 359515.1](#)

8.4.2 Checking TCP Network Protocol Buffer for Direct NFS Client

Check your TCP network buffer size to ensure that it is adequate for the speed of your servers.

By default, the network buffer size is set to 1 MB for TCP, and 2 MB for UDP. The TCP buffer size can set a limit on file transfers, which can negatively affect performance for Direct NFS Client users.

To check the current TCP buffer size, enter the following command:

```
# sysctl -a |grep -e net.ipv4.tcp_[rw]mem
```

The output of this command is similar to the following:

```
net.ipv4.tcp_rmem = 4096      87380   1048576
net.ipv4.tcp_wmem = 4096      16384   1048576
```

Oracle recommends that you set the value based on the link speed of your servers. For example, perform the following steps:

1. As root, use a text editor to open `/etc/sysctl.conf`, and add or change the following:

```
net.ipv4.tcp_rmem = 4096      87380  4194304
net.ipv4.tcp_wmem = 4096      16384  4194304
```

2. Apply your changes by running the following command:

```
# sysctl -p
```

3. Restart the network:

```
# /etc/rc.d/init.d/network restart
```

8.4.3 Creating an `orandstab` File for Direct NFS Client

Direct NFS uses a configuration file, `orandstab`, to determine the available mount points.

Create an `orandstab` file with the following attributes for each NFS server that you want to access using Direct NFS Client:

- `server`
The NFS server name.
- `local`
Up to four paths on the database host, specified by IP address or by name, as displayed using the `ifconfig` command run on the database host.
- `path`
Up to four network paths to the NFS server, specified either by IP address, or by name, as displayed using the `ifconfig` command on the NFS server.
- `export`
The exported path from the NFS server.
- `mount`
The corresponding local mount point for the exported volume.
- `mnt_timeout`
Specifies (in seconds) the time Direct NFS Client should wait for a successful mount before timing out. This parameter is optional. The default timeout is 10 minutes (600).
- `nfs_version`
Specifies the NFS protocol version used by Direct NFS Client. Possible values are NFSv3, NFSv4, NFSv4.1, and pNFS. The default version is NFSv3. If you select NFSv4.x, then you must configure the value in `orandstab` for `nfs_version`. Specify `nfs_version` as `pNFS`, if you want to use Direct NFS with Parallel NFS.
- `security_default`
Specifies the default security mode applicable for all the exported NFS server paths for a server entry. This parameter is optional. `sys` is the default value. See

the description of the `security` parameter for the supported security levels for the `security_default` parameter.

- `security`

Specifies the security level, to enable security using Kerberos authentication protocol with Direct NFS Client. This optional parameter can be specified per export-mount pair. The supported security levels for the `security_default` and `security` parameters are:

`sys`: UNIX level security AUTH_UNIX authentication based on user identifier (UID) and group identifier (GID) values. This is the default value for security parameters.

`krb5`: Direct NFS runs with plain Kerberos authentication. Server is authenticated as the real server which it claims to be.

`krb5i`: Direct NFS runs with Kerberos authentication and NFS integrity. Server is authenticated and each of the message transfers is checked for integrity.

`krb5p`: Direct NFS runs with Kerberos authentication and NFS privacy. Server is authenticated, and all data is completely encrypted.

The `security` parameter, if specified, takes precedence over the `security_default` parameter. If neither of these parameters are specified, then `sys` is the default authentication.

For NFS server Kerberos security setup, review the relevant NFS server documentation. For Kerberos client setup, review the relevant operating system documentation.

- `dontroute`

Specifies that outgoing messages should not be routed by the operating system, but instead sent using the IP address to which they are bound.

Note:

The `dontroute` option is a POSIX option, which sometimes does not work on Linux systems with multiple paths in the same subnet.

- `management`

Enables Direct NFS Client to use the management interface for SNMP queries. You can use this parameter if SNMP is running on separate management interfaces on the NFS server. The default value is the server parameter value.

- `community`

Specifies the community string for use in SNMP queries. Default value is `public`.

The following examples show three possible NFS server entries in `oranfstab`. A single `oranfstab` can have multiple NFS server entries.

Example 8-1 Using Local and Path NFS Server Entries

The following example uses both local and path. Because they are in different subnets, you do not have to specify `dontroute`.

```
server: MyDataServer1
local: 192.0.2.0
path: 192.0.2.1
```

```

local: 192.0.100.0
path: 192.0.100.1
export: /vol/oradata1 mount: /mnt/oradata1

```

Example 8-2 Using Local and Path in the Same Subnet, with dontroute

Local and path in the same subnet, where dontroute is specified:

```

server: MyDataServer2
local: 192.0.2.0
path: 192.0.2.128
local: 192.0.2.1
path: 192.0.2.129
dontroute
export: /vol/oradata2 mount: /mnt/oradata2

```

Example 8-3 Using Names in Place of IP Addresses, with Multiple Exports, management and community

```

server: MyDataServer3
local: LocalPath1
path: NfsPath1
local: LocalPath2
path: NfsPath2
local: LocalPath3
path: NfsPath3
local: LocalPath4
path: NfsPath4
dontroute
export: /vol/oradata3 mount: /mnt/oradata3
export: /vol/oradata4 mount: /mnt/oradata4
export: /vol/oradata5 mount: /mnt/oradata5
export: /vol/oradata6 mount: /mnt/oradata6
management: MgmtPath1
community: private

```

Example 8-4 Using Kerberos Authentication with Direct NFS Export

The security parameter overrides security_default:

```

server: nfsserver
local: 192.0.2.0
path: 192.0.2.2
local: 192.0.2.3
path: 192.0.2.4
export: /private/oracle1/logs mount: /logs security: krb5
export: /private/oracle1/data mount: /data security: krb5p
export: /private/oracle1/archive mount: /archive security: sys
export: /private/oracle1/data1 mount: /data1
security_default: krb5i

```

8.4.4 Enabling and Disabling Direct NFS Client Control of NFS

Use these commands to enable or disable Direct NFS Client Oracle Disk Manager Control of NFS.

By default, Direct NFS Client is installed in an enabled state. However, if Direct NFS Client is disabled and you want to enable it, complete the following steps on each node. If you use a shared Grid home for the cluster, then complete the following steps in the shared Grid home:

1. Log in as the Oracle Grid Infrastructure installation owner.

2. Change directory to *Grid_home/rdbms/lib*.
3. Enter the following command:

```
$ make -f ins_rdbms.mk dnfs_on
```

Note:

If you remove an NFS path that an Oracle Database is using, then you must restart the database for the change to take effect.

8.4.5 Enabling Hybrid Columnar Compression on Direct NFS Client

Perform these steps to enable Hybrid Columnar Compression (HCC) on Direct NFS Client:

1. Ensure that SNMP is enabled on the ZFS storage server. For example:


```
$ snmpget -v1 -c public server_name .1.3.6.1.4.1.42.2.225.1.4.2.0
SNMPv2-SMI::enterprises.42.2.225.1.4.2.0 = STRING: "Sun Storage 7410"
```
2. If SNMP is enabled on an interface other than the NFS server, then configure `oranfstab` using the management parameter.
3. If SNMP is configured using a community string other than `public`, then configure `oranfstab` file using the community parameter.
4. Ensure that `libnetsnmp.so` is installed by checking if `snmpget` is available.

8.5 Creating Member Cluster Manifest File for Oracle Member Clusters

Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Rapid Home Provisioning configuration.

Oracle Member Clusters use Oracle ASM storage from the Oracle Domain Services Cluster. Grid Naming Service (GNS) without zone delegation must be configured so that the GNS virtual IP address (VIP) is available for connection.

1. (Optional) If the Oracle Member Cluster accesses direct or indirect Oracle ASM storage, then, enable access to the disk group. Connect to any Oracle ASM instance as SYSASM user and run the command:

```
ALTER DISKGROUP diskgroup_name SET ATTRIBUTE 'access_control.enabled' = 'true';
```

2. From the Grid home on the Oracle Domain Services Cluster, create the member cluster manifest file:

```
cd Grid_home/bin
./crsctl create member_cluster_configuration member_cluster_name
-file cluster_manifest_file_name -member_type database/application [-version
member_cluster_version
[-domain_services [asm_storage local/direct/indirect][rhpf]]
```

member_cluster_name is the client cluster name.

`-file` specifies the full path of the xml file to export the credentials, `-version` is the five digit Client Cluster version, for example, `12.2.0.1.0`, if it is different

from the Storage Server version. The Storage Server version is used if `-version` is not specified.

In the options for `-domain_services`, specifying `rhp` generates credentials and configuration for a RHP Client Cluster, and `asm_storage` generates credentials and configuration for an Oracle ASM Client Cluster. `direct` if specified, signifies direct storage access, otherwise indirect.

This command creates a member cluster manifest file containing configuration details about Grid Infrastructure Management Repository (GIMR), Storage services, and Rapid Home Provisioning for the Oracle Member Cluster.

3. GNS client data is required if the Oracle Member Cluster uses dynamic networks and the server cluster has GNS with zone delegation. Provide the GNS client data as follows:
 - a. As `root` or `grid` user, export the Grid Naming Service (GNS) client data, to the member cluster manifest file created earlier:

```
srvctl export gns -clientdata manifest_file_name -role CLIENT
```

The GNS configuration is appended to the member cluster manifest file.

4. Copy the manifest file to a location on the Oracle Member Cluster, and select the file during the installation and configuration of the Oracle Member Cluster.

Related Topics:

[Installing Oracle Member Clusters](#) (page 9-21)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

8.6 Configuring Oracle Automatic Storage Management Cluster File System

Review this information to configure Oracle ACFS for an Oracle RAC Oracle Database home.

Oracle ACFS is installed as part of an Oracle Grid Infrastructure installation 12c release 2 (12.2).

You can also create a General Purpose File System configuration of ACFS using ASMCA.

To configure Oracle ACFS for an Oracle Database home for an Oracle RAC database:

1. Install Oracle Grid Infrastructure for a cluster.
2. Change directory to the Oracle Grid Infrastructure home. For example:

```
$ cd /u01/app/12.2.0/grid
```
3. Ensure that the Oracle Grid Infrastructure installation owner has read and write permissions on the storage mountpoint you want to use. For example, if you want to use the mountpoint `/u02/acfsmounts/`:

```
$ ls -l /u02/acfsmounts
```
4. Start Oracle ASM Configuration Assistant as the grid installation owner. For example:

```
./asmca
```

5. The Configure ASM: ASM Disk Groups page shows you the Oracle ASM disk group you created during installation. Click the **ASM Cluster File Systems** tab.
6. On the ASM Cluster File Systems page, right-click the Data disk, then select **Create ACFS for Database Use**.
7. In the Create ACFS for Database window, enter the following information:
 - **Volume Name:** Enter the name of the database home. The name must be unique in your enterprise. For example: `dbase_01`
 - **Mount Point:** Enter the directory path for the mount point. For example: `/u02/acfsmounts/dbase_01`
Make a note of this mount point for future reference.
 - **Size (GB):** Enter in gigabytes the size you want the database home to be. The default is 12 GB and the minimum recommended size.
 - **Owner Name:** Enter the name of the Oracle Database installation owner you plan to use to install the database. For example: `oracle1`
 - **Owner Group:** Enter the OSDBA group whose members you plan to provide when you install the database. Members of this group are given operating system authentication for the SYSDBA privileges on the database. For example: `dba1`

Select **Automatically run configuration commands** to run ASMCA configuration commands automatically. To use this option, you must provide the root credentials on the ASMCA Settings page.

Click **OK** when you have completed your entries.

8. If you did not select to run configuration commands automatically, then run the script generated by Oracle ASM Configuration Assistant as a privileged user (`root`). On an Oracle Clusterware environment, the script registers the ACFS as a resource managed by Oracle Clusterware. Registering ACFS as a resource helps Oracle Clusterware to mount ACFS automatically in proper order when ACFS is used for an Oracle RAC Oracle Database home.
9. During Oracle RAC installation, ensure that you or the DBA who installs Oracle RAC selects for the Oracle home the mount point you provided in the **Mount Point** field (in the preceding example, `/u02/acfsmounts/dbase_01`).

See Also:

Oracle Automatic Storage Management Administrator's Guide for more information about configuring and managing your storage with Oracle ACFS

8.7 Checking OCFS2 Version Manually

Use the following command to check your OCFS2 version manually.

Ensure that `ocfs2console` and `ocfs2-tools` are at least version 1.2.7, and that the other OCFS2 components correspond to the pattern `ocfs2-`

kernel_version-1.2.7 or greater. If you want to install Oracle RAC on a shared home, then the OCFS2 version must be 1.4.1 or greater.

```
modinfo ocfs2  
rpm -qa |grep ocfs2
```

For information about OCFS2, refer to the following website:

<http://oss.oracle.com/projects/ocfs2/>

Installing Oracle Grid Infrastructure

Review this information for installation and deployment options for Oracle Grid Infrastructure.

Oracle Database and Oracle Grid Infrastructure installation software is available in multiple media, and can be installed using several options. The Oracle Grid Infrastructure software is available as an image, available for download from the Oracle Technology Network website, or the Oracle Software Delivery Cloud portal. In most cases, you use the graphical user interface (GUI) provided by Oracle Universal Installer to install the software. You can also use Oracle Universal Installer to complete silent mode installations, without using the GUI. You can also use Rapid Home Provisioning for subsequent Oracle Grid Infrastructure and Oracle Database deployments.

[About Image-Based Oracle Grid Infrastructure Installation](#) (page 9-2)

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.

[Understanding Cluster Configuration Options](#) (page 9-3)

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 12c Release 2.

[Installing Oracle Grid Infrastructure for a New Cluster](#) (page 9-7)

Review these procedures to install the cluster configuration options available in this release of Oracle Grid Infrastructure.

[Installing Oracle Grid Infrastructure Using a Cluster Configuration File](#) (page 9-26)

During installation of Oracle Grid Infrastructure, you have the option of either of providing cluster configuration information manually, or of using a cluster configuration file.

[Installing Only the Oracle Grid Infrastructure Software](#) (page 9-28)

This installation option requires manual postinstallation steps to enable the Oracle Grid Infrastructure software.

[About Deploying Oracle Grid Infrastructure Using Rapid Home Provisioning](#) (page 9-31)

Rapid Home Provisioning is a software lifecycle management method for provisioning and patching Oracle homes. Rapid Home Provisioning enables mass deployment of standard operating environments for databases and clusters.

[Confirming Oracle Clusterware Function](#) (page 9-32)

After Oracle Grid Infrastructure installation, confirm that your Oracle Clusterware installation is installed and running correctly.

[Confirming Oracle ASM Function for Oracle Clusterware Files](#) (page 9-32)
Confirm Oracle ASM is running after installing Oracle Grid Infrastructure.

[Understanding Offline Processes in Oracle Grid Infrastructure](#) (page 9-33)
After the installation of Oracle Grid Infrastructure, some components may be listed as OFFLINE. Oracle Grid Infrastructure activates these resources when you choose to add them.

9.1 About Image-Based Oracle Grid Infrastructure Installation

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.

To install Oracle Grid Infrastructure, create the new Grid home with the necessary user group permissions, and then extract the image file into the newly-created Grid home, and run the setup wizard to register the Oracle Grid Infrastructure product.

Using image-based installation, you can do the following:

- Install and upgrade Oracle Grid Infrastructure for cluster configurations.
- Install Oracle Grid Infrastructure for a standalone server (Oracle Restart).
- Install only Oracle Grid Infrastructure software, and register the software with Oracle inventory.
- Add nodes to your existing cluster, if the Oracle Grid Infrastructure software is already installed or configured.

This installation feature streamlines the installation process and supports automation of large-scale custom deployments. You can also use this installation method for deployment of customized images, after you patch the base-release software with the necessary Patch Set Updates (PSUs) and patches.

Note: You must extract the image software into the directory where you want your Grid home to be located, and then run the `gridSetup.sh` script to start the Grid Infrastructure setup wizard. Ensure that the Grid home directory path you create is in compliance with the Oracle Optimal Flexible Architecture recommendations.

Related Topics:

[Installing Oracle Domain Services Cluster](#) (page 9-14)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.

[Installing Oracle Standalone Cluster](#) (page 9-7)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

[Installing Oracle Member Clusters](#) (page 9-21)
Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

9.2 Understanding Cluster Configuration Options

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 12c Release 2.

[About Oracle Standalone Clusters](#) (page 9-3)

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

[About Oracle Cluster Domain and Oracle Domain Services Cluster](#) (page 9-3)

An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.

[About Oracle Member Clusters](#) (page 9-5)

Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications.

[About Oracle Extended Clusters](#) (page 9-6)

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

9.2.1 About Oracle Standalone Clusters

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

Oracle Standalone Clusters contain two types of nodes arranged in a hub and spoke architecture: Hub Nodes and Leaf Nodes. The number of Hub Nodes in an Oracle Standalone Cluster can be as many as 64. The number of Leaf Nodes can be many more. Hub Nodes and Leaf Nodes can host different types of applications. Oracle Standalone Cluster Hub Nodes are tightly connected, and have direct access to shared storage. Leaf Nodes do not require direct access to shared storage. Hub Nodes can run in an Oracle Standalone Cluster configuration without having any Leaf Nodes as cluster member nodes, but Leaf Nodes must be members of a cluster with a pool of Hub Nodes. Shared storage is locally mounted on each of the Hub Nodes, with an Oracle ASM instance available to all Hub Nodes.

Oracle Standalone Clusters host Grid Infrastructure Management Repository (GIMR) locally. The GIMR is a multitenant database, which stores information about the cluster. This information includes the real time performance data the Cluster Health Monitor collects, and includes metadata required for Rapid Home Provisioning.

When you deploy an Oracle Standalone Cluster, you can also choose to configure it as an Oracle Extended cluster. An Oracle Extended Cluster consists of nodes that are located in multiple locations or sites.

9.2.2 About Oracle Cluster Domain and Oracle Domain Services Cluster

An Oracle Cluster Domain is a choice of deployment architecture for new clusters, introduced in Oracle Clusterware 12c Release 2.

Oracle Cluster Domain enables you to standardize, centralize, and optimize your Oracle Real Application Clusters (Oracle RAC) deployment for the private database cloud. Multiple cluster configurations are grouped under an Oracle Cluster Domain for management purposes and make use of shared services available within that

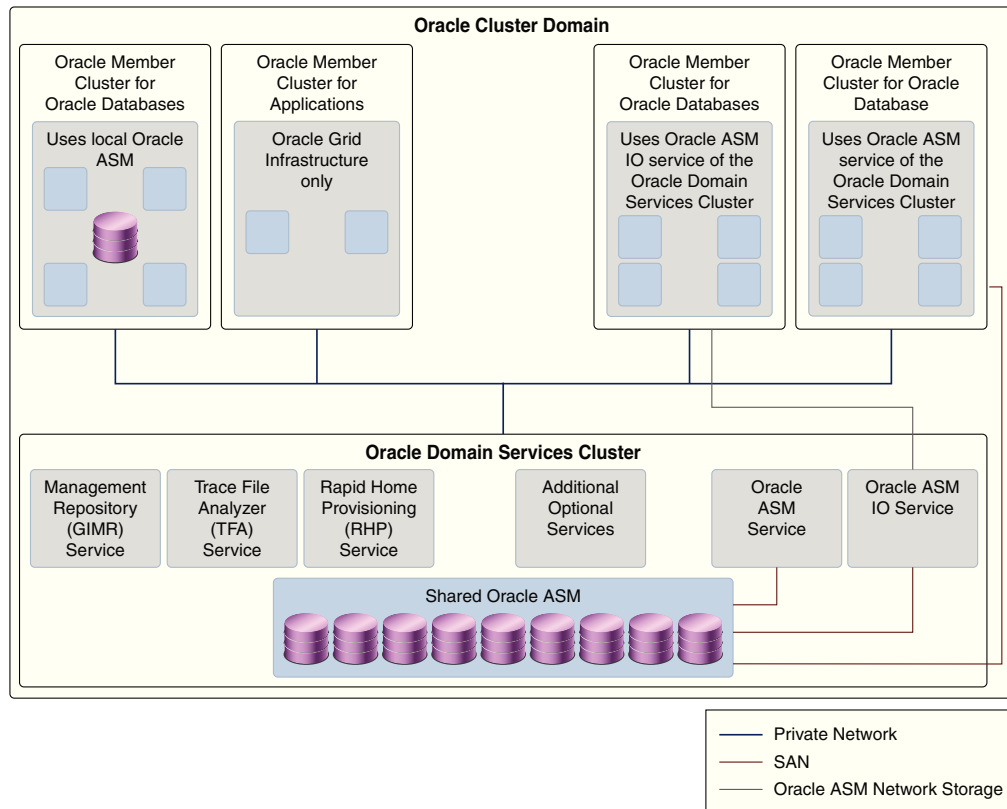
Oracle Cluster Domain. The cluster configurations within that Oracle Cluster Domain include Oracle Domain Services Cluster and Oracle Member Clusters.

The Oracle Domain Services Cluster provides centralized services to other clusters within the Oracle Cluster Domain. These services include:

- A centralized Grid Infrastructure Management Repository (housing the MGMTDB for each of the clusters within the Oracle Cluster Domain)
- Trace File Analyzer (TFA) services, for targeted diagnostic data collection for Oracle Clusterware and Oracle Database
- Consolidated Oracle ASM storage management service
- An optional Rapid Home Provisioning (RHP) Service to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes. When you configure the Oracle Domain Services Cluster, you can also choose to configure the Rapid Home Provisioning Server.

An Oracle Domain Services Cluster provides these centralized services to Oracle Member Clusters. Oracle Member Clusters use these services for centralized management and to reduce their local resource usage.

Figure 9-1 Oracle Cluster Domain



Related Topics:

[About Oracle Member Clusters](#) (page 9-5)

Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications.

9.2.3 About Oracle Member Clusters

Oracle Member Clusters use centralized services from the Oracle Domain Services Cluster and can host databases or applications.

Oracle Member Clusters can be of two types — Oracle Member Clusters for Oracle Databases or Oracle Member Clusters for applications.

Oracle Member Clusters do not need direct connectivity to shared disks. Using the shared Oracle ASM service, they can use network connectivity to the IO Service to access a centrally managed pool of storage. To use shared Oracle ASM services from the Oracle Domain Services Cluster, the member cluster needs connectivity to the Oracle ASM networks of the Oracle Domain Services Cluster.

Oracle Member Clusters cannot provide services to other clusters. For example, you cannot configure and use a member cluster as a GNS server or Rapid Home Provisioning Server.

Oracle Member Cluster for Oracle Databases

An Oracle Member Cluster for Oracle Databases supports Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database instances. This cluster registers with the management repository service and uses the centralized TFA service. It can use additional services as needed. An Oracle Member Cluster for Oracle Databases can be configured with local Oracle ASM storage management or make use of the consolidated Oracle ASM storage management service offered by the Oracle Domain Services Cluster.

An Oracle Member Cluster for Oracle Database always uses remote Grid Infrastructure Management Repository (GIMR) from its Oracle Domain Services Cluster. For two-node or four-node clusters, hosting the GIMR on a remote cluster reduces the overhead of running an extra infrastructure repository on a cluster.

Oracle Member Cluster for Applications

Oracle Member Cluster for Applications hosts applications other than Oracle Database, as part of an Oracle Cluster Domain. The Oracle Member Cluster requires connectivity to Oracle Cluster Domain Services for centralized management and resource efficiency. The Oracle Member Cluster uses remote Oracle ASM storage and does not require direct shared storage access. This cluster configuration enables high availability of any software application.

Note:

Before running Oracle Universal Installer, you must specify the Oracle Domain Services Cluster configuration details for the Oracle Member Cluster by creating the Member Cluster Manifest file.

Oracle Member Cluster for Oracle Database does not support Oracle Database 12.1 or earlier, where Oracle Member Cluster is configured with Oracle ASM storage as direct or indirect.

Related Topics:

[Creating Member Cluster Manifest File for Oracle Member Clusters](#) (page 8-21)
Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management

Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Rapid Home Provisioning configuration.

9.2.4 About Oracle Extended Clusters

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

When you deploy an Oracle Standalone Cluster, you can also choose to configure the cluster as an Oracle Extended Cluster. You can extend an Oracle RAC cluster across two, or more, geographically separate sites, each equipped with its own storage. In the event that one of the sites fails, the other site acts as an active standby.

Both Oracle ASM and the Oracle Database stack, in general, are designed to use enterprise-class shared storage in a data center. Fibre Channel technology, however, enables you to distribute compute and storage resources across two or more data centers, and connect them through Ethernet cables and Fibre Channel, for compute and storage needs, respectively.

You can configure an Oracle Extended Cluster when you install Oracle Grid Infrastructure. You can also do so post installation using the `ConvertToExtended` script. You manage your Oracle Extended Cluster using `CRSCTL`.

Oracle recommends that you deploy Oracle Extended Clusters with normal redundancy disk groups. You can assign nodes and failure groups to sites. Sites contain failure groups, and failure groups contain disks. For normal redundancy disk groups, a disk group provides one level of failure protection, and can tolerate the failure of either a site or a failure group.

The following conditions apply when you select redundancy levels for Oracle Extended Clusters:

Table 9-1 Oracle ASM Disk Group Redundancy Levels for Oracle Extended Clusters

Redundancy Level	Number of OCR and Voting Files Disk Groups	Number of OCR Backup and GIMR Disk Groups
Normal redundancy	1 failure group per data site, 1 quorum failure group	1 failure group per data site
Flex redundancy	1 failure group per data site, 1 quorum failure group	Three failure groups, with 1 failure group per site
High redundancy	Not supported	Three failure groups, with 1 failure group per site

Related Topics:

[Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure](#) (page 11-32)

Review this information to convert to an Oracle Extended Cluster after upgrading Oracle Grid Infrastructure. Oracle Extended Cluster enables you to deploy Oracle RAC databases on a cluster, in which some of the nodes are located in different sites.

See Also: *Oracle Clusterware Administration and Deployment Guide*

9.3 Installing Oracle Grid Infrastructure for a New Cluster

Review these procedures to install the cluster configuration options available in this release of Oracle Grid Infrastructure.

[Installing Oracle Standalone Cluster](#) (page 9-7)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

[Installing Oracle Domain Services Cluster](#) (page 9-14)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.

[Installing Oracle Member Clusters](#) (page 9-21)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

9.3.1 Installing Oracle Standalone Cluster

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), the installation media is replaced with a zip file for the Oracle Grid Infrastructure installer. Run the installation wizard after extracting the zip file into the target home path.

At any time during installation, if you have a question about what you are being asked to do, or what input you are required to provide during installation, click the **Help** button on the installer page.

You should have your network information, storage information, and operating system users and groups available to you before you start installation, and you should be prepared to run root scripts.

As the user that owns the software for Oracle Grid Infrastructure for a cluster (`grid`) on the first node, install Oracle Grid Infrastructure for a cluster. Note that the installer uses Secure Shell (SSH) to copy the binary files from this node to the other nodes during the installation. During installation, in the Cluster Node Information window, when you specify the nodes in your cluster, you can click **SSH Connectivity** and the installer configures SSH connectivity between the specified nodes for you.

Note: These installation instructions assume you do not already have any Oracle software installed on your system. If you have already installed Oracle ASMLIB, then you cannot install Oracle ASM Filter Driver (Oracle ASMPFD) until you uninstall Oracle ASMLIB. You can use Oracle ASMLIB instead of Oracle ASMPFD for managing the disks used by Oracle ASM.

To install the software for Oracle Standalone Cluster:

1. As the `grid` user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home. For example:

```
mkdir -p /u01/app/12.2.0/grid
chown grid:oinstall /u01/app/12.2.0/grid
cd /u01/app/12.2.0/grid
unzip -q download_location/grid.zip
```

`grid.zip` is the name of the Oracle Grid Infrastructure image zip file. For example, on Linux systems, the name of the Oracle Grid Infrastructure image zip file is `linuxx64_12201_grid_home.zip`.

Note:

- You must extract the zip image software into the directory where you want your Grid home to be located.
 - Download and copy the Oracle Grid Infrastructure image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.
-

2. Configure the shared disks for use with Oracle ASM Filter Driver:

- a. Verify that the device rules file to set the permissions and ownership is created at `/etc/udev/rules.d/`. See *Configuring Device Persistence Manually for Oracle ASM* for the procedure to create the device rules file.
- b. Log in as the `root` user and set the environment variable `ORACLE_HOME` to the location of the Grid home.

For C shell:

```
su root
setenv ORACLE_HOME /u01/app/12.2.0/grid
```

For bash shell:

```
su root
export ORACLE_HOME=/u01/app/12.2.0/grid
```

- c. Use Oracle ASM command line tool (ASMCMD) to provision the disk devices for use with Oracle ASM Filter Driver.

```
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA1 /dev/sdb1 --init
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA2 /dev/sdc1 --init
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA3 /dev/sdd1 --init
```

- d. Verify the device has been marked for use with Oracle ASMFDD.

```
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdb1
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdc1
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdd1
```

3. Log in as the `grid` user, and start the Oracle Grid Infrastructure installer by running the following command:

```
Grid_home/gridSetup.sh
```

The installer starts and the Select Configuration Option window appears.

4. Choose the option **Configure Grid Infrastructure for a New Cluster**, then click **Next**.

The Select Cluster Configuration window appears.

5. Choose the option **Configure an Oracle Standalone Cluster**, then click **Next**.

Select the **Configure as Extended Cluster** option to extend an Oracle RAC cluster across two or more separate sites, each equipped with its own storage.

The Grid Plug and Play Information window appears.

6. In the **Cluster Name** and **SCAN Name** fields, enter the names for your cluster and cluster scan that are unique throughout your entire enterprise network.

You can select **Configure GNS** if you have configured your domain name server (DNS) to send to the GNS virtual IP address name resolution requests for the subdomain GNS serves, as explained in this guide.

For cluster member node public and VIP network addresses, provide the information required depending on the kind of cluster you are configuring:

- If you plan to use automatic cluster configuration with DHCP addresses configured and resolved through GNS, then you only need to provide the GNS VIP names as configured on your DNS.
- If you plan to use manual cluster configuration, with fixed IP addresses configured and resolved on your DNS, then provide the SCAN names for the cluster, and the public names, and VIP names for each cluster member node. For example, you can choose a name that is based on the node names' common prefix. The cluster name can be `mycluster` and the cluster SCAN name can be `mycluster-scan`.

Click **Next**.

The Cluster Node Information screen appears.

7. In the Public Hostname column of the table of cluster nodes, you should see your local node, for example `node1.example.com`.

The following is a list of additional information about node IP addresses:

- For the local node only, OUI automatically fills in public and VIP fields. If your system uses vendor clusterware, then OUI may fill additional fields.
 - Host names and virtual host names are not domain-qualified. If you provide a domain in the address field during installation, then OUI removes the domain from the address.
 - Interfaces identified as private for private IP addresses should not be accessible as public interfaces. Using public interfaces for Cache Fusion can cause performance problems.
 - When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `/bin/hostname` command.
- a. Click **Add** to add another node to the cluster.
 - b. Enter the second node's public name (`node2`), and virtual IP name (`node2-vip`), then click **OK**.

You are returned to the Cluster Node Information window. You should now see all nodes listed in the table of cluster nodes. Make sure the Role column is set to HUB for both nodes. To add Leaf Nodes, you must configure GNS.

- c. Make sure all nodes are selected, then click the **SSH Connectivity** button at the bottom of the window.

The bottom panel of the window displays the SSH Connectivity information.

- d. Enter the operating system user name and password for the Oracle software owner (`grid`). Select the option **If you have configured SSH connectivity between the nodes**, then select the **Reuse private and public keys existing in user home** option. Click **Setup**.

A message window appears, indicating that it might take several minutes to configure SSH connectivity between the nodes. After a short period, another message window appears indicating that passwordless SSH connectivity has been established between the cluster nodes. Click **OK** to continue.

- e. When returned to the Cluster Node Information window, click **Next** to continue.

The Specify Network Interface Usage page appears.

8. Select the usage type for each network interface displayed.

Verify that each interface has the correct interface type associated with it. If you have network interfaces that should not be used by Oracle Clusterware, then set the network interface type to **Do Not Use**. For example, if you have only two network interfaces, then set the public interface to have a Use For value of **Public** and set the private network interface to have a Use For value of **ASM & Private**.

Click **Next**. The Storage Option Information window appears.

9. Select the Oracle ASM storage configuration option:
 - a. If you select **Configure ASM using block devices**, then specify the NFS mount points for the Oracle ASM disk groups, and optionally, the GIMR disk group in the Specify NFS Locations for ASM Disk Groups window.
 - b. If you select **Configure ASM on NAS**, then click **Next**. The Grid Infrastructure Management Repository Option window appears.
10. Choose whether you want to store the Grid Infrastructure Management Repository in a separate Oracle ASM disk group, then click **Next**.

The Create ASM Disk Group window appears.

11. Provide the name and specifications for the Oracle ASM disk group.
 - a. In the **Disk Group Name** field, enter a name for the disk group, for example `DATA`.
 - b. Choose the Redundancy level for this disk group. Normal is the recommended option.
 - c. In the Add Disks section, choose the disks to add to this disk group.

In the Add Disks section you should see the disks that you labeled in Step 2. If you do not see the disks, click the **Change Discovery Path** button and provide a path and pattern match for the disk, for example, `/dev/sd*`

During installation, disks labelled as Oracle ASMFD disks or Oracle ASMLIB disks are listed as candidate disks when using the default discovery string. However, if the disk has a header status of MEMBER, then it is not a candidate disk.

- d. If you want to use Oracle ASM Filter Driver (Oracle ASMFD) to manage your Oracle ASM disk devices, then select the option **Configure Oracle ASM Filter Driver**.

If you are installing on Linux systems, and you want to use Oracle ASM Filter Driver (Oracle ASMFD) to manage your Oracle ASM disk devices, then you must deinstall Oracle ASM library driver (Oracle ASMLIB) before starting Oracle Grid Infrastructure installation.

When you have finished providing the information for the disk group, click **Next**.

12. If you selected to use a different disk group for the GIMR, then the Grid Infrastructure Management Repository Option window appears. Provide the name and specifications for the GIMR disk group.

- a. In the **Disk Group Name** field, enter a name for the disk group, for example DATA.
- b. Choose the Redundancy level for this disk group. Normal is the recommended option.
- c. In the Add Disks section, choose the disks to add to this disk group.

When you have finished providing the information for the disk group, click **Next**.

The Specify ASM Password window appears.

13. Choose the same password for the Oracle ASM SYS and ASMSNMP account, or specify different passwords for each account, then click **Next**.

The Failure Isolation Support window appears.

14. Select the option **Do not use Intelligent Platform Management Interface (IPMI)**, then click **Next**.

The Specify Management Options window appears.

15. If you have Enterprise Manager Cloud Control installed in your enterprise, then choose the option **Register with Enterprise Manager (EM) Cloud Control** and provide the EM configuration information. If you do not have Enterprise Manager Cloud Control installed in your enterprise, then click **Next** to continue.

The Privileged Operating System Groups window appears.

16. Accept the default operating system group names for Oracle ASM administration and click **Next**.

The Specify Install Location window appears.

17. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation, then click **Next**. The Oracle base directory must be different from the Oracle home directory.

If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as directed in Step 1, then the default location for the Oracle base directory should display as `/u01/app/grid`.

If you have not installed Oracle software previously on this computer, then the Create Inventory window appears.

18. Change the path for the inventory directory, if required. Then, click **Next**.

If you are using the same directory names as the examples in this book, then it should show a value of `/u01/app/oraInventory`. The group name for the `oraInventory` directory should show `oinstall`.

The Root Script Execution Configuration window appears.

19. Select the option to **Automatically run configuration scripts**. Enter the credentials for the root user or a sudo account, then click **Next**.

Alternatively, you can Run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

The Perform Prerequisite Checks window appears.

20. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the **Check Again** button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

The Summary window appears.

21. Review the contents of the Summary window and then click **Install**.

The installer displays a progress indicator enabling you to monitor the installation process.

22. If you did not configure automation of the root scripts, then you are required to run certain scripts as the `root` user, as specified in the Execute Configuration Scripts window appears. Do *not* click **OK** until you have run the scripts. Run the scripts on all nodes as directed, in the order shown.

For example, on Oracle Linux you perform the following steps (note that for clarity, the examples show the current user, node and directory in the prompt):

- a. As the `oracle` user on `node1`, open a terminal window, and enter the following commands:

```
[oracle@node1 oracle]$ cd /u01/app/oraInventory
[oracle@node1 oraInventory]$ su
```

- b. Enter the password for the `root` user, and then enter the following command to run the first script on `node1`:

```
[root@node1 oraInventory]# ./orainstRoot.sh
```

- c. After the `orainstRoot.sh` script finishes on `node1`, open another terminal window, and as the `oracle` user, enter the following commands:

```
[oracle@node1 oracle]$ ssh node2
[oracle@node2 oracle]$ cd /u01/app/oraInventory
[oracle@node2 oraInventory]$ su
```


- d. Enter the password for the `root` user, and then enter the following command to run the first script on `node2`:

```
[root@node2 oraInventory]# ./oraInstRoot.sh
```

- e. After the `oraInstRoot.sh` script finishes on `node2`, go to the terminal window you opened in part a of this step. As the `root` user on `node1`, enter the following commands to run the second script, `root.sh`:

```
[root@node1 oraInventory]# cd /u01/app/12.2.0/grid
[root@node1 grid]# ./root.sh
```

Press **Enter** at the prompt to accept the default value.

Note:

You must run the `root.sh` script on the first node and wait for it to finish. You can run `root.sh` scripts concurrently on all other nodes except for the last node on which you run the script. Like the first node, the `root.sh` script on the last node must be run separately.

- f. After the `root.sh` script finishes on `node1`, go to the terminal window you opened in part c of this step. As the `root` user on `node2`, enter the following commands:

```
[root@node2 oraInventory]# cd /u01/app/12.2.0/grid
[root@node2 grid]# ./root.sh
```

After the `root.sh` script completes, return to the Oracle Universal Installer window where the Installer prompted you to run the `oraInstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

23. Continue monitoring the installation until the Finish window appears. Then click **Close** to complete the installation process and exit the installer.

Caution:

After installation is complete, do not remove manually or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running on the server. If you remove these files, then the Oracle software can encounter intermittent hangs. Oracle Clusterware installations can fail with the error:

```
CRS-0184: Cannot communicate with the CRS daemon.
```

After your Oracle Grid Infrastructure installation is complete, you can install Oracle Database on a cluster node for high availability, or install Oracle RAC.

See Also: *Oracle Real Application Clusters Installation Guide* or *Oracle Database Installation Guide* for your platform for information on installing Oracle Database

Related Topics:

[Configuring Device Persistence Manually for Oracle ASM](#) (page B-15)

Complete these tasks to create device path persistence manually for Oracle ASM.

9.3.2 Installing Oracle Domain Services Cluster

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), the installation media is replaced with a zip file for the Oracle Grid Infrastructure installer. Run the installation wizard after extracting the zip file into the target home path.

At any time during installation, if you have a question about what you are being asked to do, or what input you are required to provide during installation, click the **Help** button on the installer page.

You should have your network information, storage information, and operating system users and groups available to you before you start installation, and you should be prepared to run root scripts.

As the user that owns the software for Oracle Grid Infrastructure for a cluster (`grid`) on the first node, install Oracle Grid Infrastructure for a cluster. Note that the installer uses Secure Shell (SSH) to copy the binary files from this node to the other nodes during the installation. During installation, in the Cluster Node Information window, when you specify the nodes in your cluster, you can click **SSH Connectivity** and the installer configures SSH connectivity between the specified nodes for you.

Note: These installation instructions assume you do not already have any Oracle software installed on your system. If you have already installed Oracle ASMLIB, then you cannot install Oracle ASM Filter Driver (Oracle ASMFDD) until you uninstall Oracle ASMLIB. You can use Oracle ASMLIB instead of Oracle ASMFDD for managing the disks used by Oracle ASM.

To install the software for Oracle Domain Services Cluster:

1. As the `grid` user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home. For example:

```
mkdir -p /u01/app/12.2.0/grid
chown grid:oinstall /u01/app/12.2.0/grid
cd /u01/app/12.2.0/grid
unzip -q download_location/grid.zip
```

`grid.zip` is the name of the Oracle Grid Infrastructure image zip file. For example, on Linux systems, the name of the Oracle Grid Infrastructure image zip file is `linuxx64_12201_grid_home.zip`.

Note:

- You must extract the zip image software into the directory where you want your Grid home to be located.
- Download and copy the Oracle Grid Infrastructure image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.

2. Configure the shared disks for use with Oracle ASM Filter Driver:

- Verify that the device rules file to set the permissions and ownership is created at `/etc/udev/rules.d/`. See *Configuring Device Persistence Manually for Oracle ASM* for the procedure to create the device rules file.
- Log in as the `root` user and set the environment variable `ORACLE_HOME` to the location of the Grid home.

For C shell:

```
su root
setenv ORACLE_HOME /u01/app/12.2.0/grid
```

For bash shell:

```
su root
export ORACLE_HOME=/u01/app/12.2.0/grid
```

- Use Oracle ASM command line tool (ASMCMD) to provision the disk devices for use with Oracle ASM Filter Driver.

```
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA1 /dev/sdb1 --init
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA2 /dev/sdc1 --init
./u01/app/12.2.0/grid/bin/asmcmd afd_label DATA3 /dev/sdd1 --init
```

- Verify the device has been marked for use with Oracle ASMFDD.

```
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdb1
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdc1
./u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/sdd1
```

- Log in as the `grid` user, and start the Oracle Grid Infrastructure installer by running the following command:

```
Grid_home/gridSetup.sh
```

The installer starts and the Select Configuration Option window appears.

- Choose the option **Configure Grid Infrastructure for a New Cluster**, then click **Next**.

The Select Cluster Configuration window appears.

- Choose the option **Configure an Oracle Domain Services Cluster**, then click **Next**.

The Grid Plug and Play Information window appears.

- In the **Cluster Name** and **SCAN Name** fields, enter the names for your cluster and cluster scan that are unique throughout your entire enterprise network.

You can select **Configure GNS** if you have configured your domain name server (DNS) to send to the GNS virtual IP address name resolution requests for the subdomain GNS serves, as explained in this guide.

For cluster member node public and VIP network addresses, provide the information required depending on the kind of cluster you are configuring:

- If you plan to use automatic cluster configuration with DHCP addresses configured and resolved through GNS, then you only need to provide the GNS VIP names as configured on your DNS.
- If you plan to use manual cluster configuration, with fixed IP addresses configured and resolved on your DNS, then provide the SCAN names for the cluster, and the public names, and VIP names for each cluster member node. For example, you can choose a name that is based on the node names' common prefix. This example uses the cluster name `mycluster` and the cluster SCAN name of `mycluster-scan`.

Click **Next**.

The Cluster Node Information screen appears.

7. In the Public Hostname column of the table of cluster nodes, you should see your local node, for example `node1.example.com`.

The following is a list of additional information about node IP addresses:

- For the local node only, OUI automatically fills in public and VIP fields. If your system uses vendor clusterware, then OUI may fill additional fields.
- Host names and virtual host names are not domain-qualified. If you provide a domain in the address field during installation, then OUI removes the domain from the address.
- Interfaces identified as private for private IP addresses should not be accessible as public interfaces. Using public interfaces for Cache Fusion can cause performance problems.
- When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `/bin/hostname` command.

- a. Click **Add** to add another node to the cluster.
- b. Enter the second node's public name (`node2`), and virtual IP name (`node2-vip`), then click **OK**.

You are returned to the Cluster Node Information window. You should now see all nodes listed in the table of cluster nodes. Make sure the Role column is set to HUB for both nodes. To add Leaf Nodes, you must configure GNS.

- c. Make sure all nodes are selected, then click the **SSH Connectivity** button at the bottom of the window.

The bottom panel of the window displays the SSH Connectivity information.

- d. Enter the operating system user name and password for the Oracle software owner (`grid`). Select the option **If you have configured SSH connectivity**

between the nodes, then select the **Reuse private and public keys existing in user home** option. Click **Setup**.

A message window appears, indicating that it might take several minutes to configure SSH connectivity between the nodes. After a short period, another message window appears indicating that passwordless SSH connectivity has been established between the cluster nodes. Click **OK** to continue.

- e. When returned to the Cluster Node Information window, click **Next** to continue.

The Specify Network Interface Usage page appears.

8. Select the usage type for each network interface displayed.

Verify that each interface has the correct interface type associated with it. If you have network interfaces that should not be used by Oracle Clusterware, then set the network interface type to **Do Not Use**. For example, if you have only two network interfaces, then set the public interface to have a Use For value of **Public** and set the private network interface to have a Use For value of **ASM & Private**.

Click **Next**. The Create ASM Disk Group window appears.

9. Provide the name and specifications for the Oracle ASM disk group.
 - a. In the **Disk Group Name** field, enter a name for the disk group, for example DATA.
 - b. Choose the Redundancy level for this disk group. Normal is the recommended option.
 - c. In the Add Disks section, choose the disks to add to this disk group.

In the Add Disks section you should see the disks that you labeled in Step 2. If you do not see the disks, click the **Change Discovery Path** button and provide a path and pattern match for the disk, for example, `/dev/sd*`.

During installation, disks labelled as Oracle ASMFD disks or Oracle ASMLIB disks are listed as candidate disks when using the default discovery string. However, if the disk has a header status of MEMBER, then it is not a candidate disk.

- d. Check the option **Configure Oracle ASM Filter Driver**.

If you are installing on Linux systems, and you want to use Oracle ASM Filter Driver (Oracle ASMFD) to manage your Oracle ASM disk devices, then you must deinstall Oracle ASM library driver (Oracle ASMLIB) before starting Oracle Grid Infrastructure installation.

When you have finished providing the information for the disk group, click **Next**.

The Grid Infrastructure Management Repository Option window appears

10. Provide the name and specifications for the GIMR disk group.
 - a. In the **Disk Group Name** field, enter a name for the disk group, for example DATA1.
 - b. Choose the Redundancy level for this disk group. Normal is the recommended option.

- c. In the Add Disks section, choose the disks to add to this disk group.
- d. Select the **Configure Rapid Home Provisioning Server** option to configure a Rapid Home Provisioning Server as part of the Oracle Domain Services Cluster. Rapid Home Provisioning enables you to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes.

When you have finished providing the information for the disk group, click **Next**.

The Specify ASM Password window appears.

- 11. Choose the same password for the Oracle ASM SYS and ASMSNMP account, or specify different passwords for each account, then click **Next**.

The Failure Isolation Support window appears.

- 12. Select the option **Do not use Intelligent Platform Management Interface (IPMI)**, then click **Next**.

The Specify Management Options window appears.

- 13. If you have Enterprise Manager Cloud Control installed in your enterprise, then choose the option **Register with Enterprise Manager (EM) Cloud Control** and provide the EM configuration information. If you do not have Enterprise Manager Cloud Control installed in your enterprise, then click **Next** to continue.

You can manage Oracle Grid Infrastructure and Oracle Automatic Storage Management (Oracle ASM) using Oracle Enterprise Manager Cloud Control. To register the Oracle Grid Infrastructure cluster with Oracle Enterprise Manager, ensure that Oracle Management Agent is installed and running on all nodes of the cluster.

The Privileged Operating System Groups window appears.

- 14. Accept the default operating system group names for Oracle ASM administration and click **Next**.

The Specify Install Location window appears.

- 15. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation, then click **Next**. The Oracle base directory must be different from the Oracle home directory.

If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as directed in Step 1, then the default location for the Oracle base directory should display as `/u01/app/grid`.

If you have not installed Oracle software previously on this computer, then the Create Inventory window appears.

- 16. Change the path for the inventory directory, if required. Then, click **Next**.

If you are using the same directory names as the examples in this book, then it should show a value of `/u01/app/orainventory`. The group name for the `oraInventory` directory should show `oinstall`.

The Root Script Execution Configuration window appears.

- 17. Select the option to **Automatically run configuration scripts**. Enter the credentials for the root user or a sudo account, then click **Next**.

Alternatively, you can Run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

The Perform Prerequisite Checks window appears.

18. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the **Check Again** button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

The Summary window appears.

19. Review the contents of the Summary window and then click **Install**.

The installer displays a progress indicator enabling you to monitor the installation process.

20. If you did not configure automation of the root scripts, then you are required to run certain scripts as the `root` user, as specified in the Execute Configuration Scripts window appears. Do *not* click **OK** until you have run the scripts. Run the scripts on all nodes as directed, in the order shown.

For example, on Oracle Linux you perform the following steps (note that for clarity, the examples show the current user, node and directory in the prompt):

- a. As the `oracle` user on `node1`, open a terminal window, and enter the following commands:

```
[oracle@node1 oracle]$ cd /u01/app/oraInventory
[oracle@node1 oraInventory]$ su
```

- b. Enter the password for the `root` user, and then enter the following command to run the first script on `node1`:

```
[root@node1 oraInventory]# ./oraInstRoot.sh
```

- c. After the `oraInstRoot.sh` script finishes on `node1`, open another terminal window, and as the `oracle` user, enter the following commands:

```
[oracle@node1 oracle]$ ssh node2
[oracle@node2 oracle]$ cd /u01/app/oraInventory
[oracle@node2 oraInventory]$ su
```

- d. Enter the password for the `root` user, and then enter the following command to run the first script on `node2`:

```
[root@node2 oraInventory]# ./oraInstRoot.sh
```

- e. After the `oraInstRoot.sh` script finishes on `node2`, go to the terminal window you opened in part a of this step. As the `root` user on `node1`, enter the following commands to run the second script, `root.sh`:

```
[root@node1 oraInventory]# cd /u01/app/12.2.0/grid
[root@node1 grid]# ./root.sh
```

Press **Enter** at the prompt to accept the default value.

Note:

You must run the `root.sh` script on the first node and wait for it to finish. If your cluster has three or more nodes, then `root.sh` can be run concurrently on all nodes but the first. Node numbers are assigned according to the order of running `root.sh`. If you want to create a particular node number assignment, then run the `root` scripts in the order of the node assignments you want to make, and wait for the script to finish running on each node before proceeding to run the script on the next node. However, Oracle system identifier, or SID, for your Oracle RAC databases, do not follow the node numbers.

- f. After the `root.sh` script finishes on `node1`, go to the terminal window you opened in part c of this step. As the `root` user on `node2`, enter the following commands:

```
[root@node2 oraInventory]# cd /u01/app/12.2.0/grid
[root@node2 grid]# ./root.sh
```

After the `root.sh` script completes, return to the OUI window where the Installer prompted you to run the `oraInstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

When you run `root.sh` during Oracle Grid Infrastructure installation, the Trace File Analyzer (TFA) Collector is also installed in the directory `grid_home/tfa`.

21. After `root.sh` runs on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
22. During the installation, Oracle Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM for storage.
23. Continue monitoring the installation until the Finish window appears. Then click **Close** to complete the installation process and exit the installer.

Caution:

After installation is complete, do not remove manually or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running on the server. If you remove these files, then the Oracle software can encounter intermittent hangs. Oracle Clusterware installations can fail with the error:

```
CRS-0184: Cannot communicate with the CRS daemon.
```

After your Oracle Domain Services Cluster installation is complete, you can install Oracle Member Clusters for Oracle Databases and Oracle Member Clusters for Applications.

Related Topics:

[Configuring Device Persistence Manually for Oracle ASM](#) (page B-15)

Complete these tasks to create device path persistence manually for Oracle ASM.

9.3.3 Installing Oracle Member Clusters

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), the installation media is replaced with a zip file for the Oracle Grid Infrastructure installer. Run the installation wizard after extracting the zip file into the target home path.

At any time during installation, if you have a question about what you are being asked to do, or what input you are required to provide during installation, click the **Help** button on the installer page.

You should have your network information, storage information, and operating system users and groups available to you before you start installation, and you should be prepared to run root scripts. Ensure that you have created a Member Cluster Manifest File as explained in this guide.

As the user that owns the software for Oracle Grid Infrastructure for a cluster (`grid`) on the first node, install Oracle Grid Infrastructure for a cluster. Note that the installer uses Secure Shell (SSH) to copy the binary files from this node to the other nodes during the installation. During installation, in the Cluster Node Information window, when you specify the nodes in your cluster, you can click **SSH Connectivity** and the installer configures SSH connectivity between the specified nodes for you.

Note: These installation instructions assume you do not already have any Oracle software installed on your system. If you have already installed Oracle ASMLIB, then you cannot install Oracle ASM Filter Driver (Oracle ASMPD) until you uninstall Oracle ASMLIB. You can use Oracle ASMLIB instead of Oracle ASMPD for managing the disks used by Oracle ASM.

To install the software for Oracle Member Cluster for Oracle Databases or Applications:

Create a Member Cluster Manifest File as explained in this guide.

Use this procedure to install an Oracle Member Cluster for Oracle Databases or Oracle Member Cluster for Applications.

1. As the `grid` user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home. For example:

```
mkdir -p /u01/app/12.2.0/grid
chown grid:oinstall /u01/app/12.2.0/grid
cd /u01/app/12.2.0/grid
unzip -q download_location/grid.zip
```

`grid.zip` is the name of the Oracle Grid Infrastructure image zip file. For example, on Linux systems, the name of the Oracle Grid Infrastructure image zip file is `linuxx64_12201_grid_home.zip`.

Note:

- You must extract the zip image software into the directory where you want your Grid home to be located.
 - Download and copy the Oracle Grid Infrastructure image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.
-
-

2. Log in as the `grid` user, and start the Oracle Grid Infrastructure installer by running the following command:

```
Grid_home/gridSetup.sh
```

The installer starts and the Select Configuration Option window appears.

3. Choose the option **Configure Grid Infrastructure for a New Cluster**, then click **Next**.

The Select Cluster Configuration window appears.

4. Choose either the **Configure an Oracle Member Cluster for Oracle Databases** or **Configure an Oracle Member Cluster for Applications** option, then click **Next**.

The Cluster Domain Services window appears.

5. Select the Manifest file that contains the configuration details about the management repository and other services for the Oracle Member Cluster.

For Oracle Member Cluster for Oracle Databases, you can also specify the Grid Naming Service and Oracle ASM Storage server details using a Member Cluster Manifest file.

Click **Next**.

6. If you selected to configure an Oracle Member Cluster for applications, then the Configure Virtual Access window appears. Provide a **Cluster Name** and optional **Virtual Host Name**.

The virtual host name serves as a connection address for the Oracle Member Cluster, and to provide service access to the software applications that you want the Oracle Member Cluster to install and run.

Click **Next**.

The Cluster Node Information screen appears.

7. In the Public Hostname column of the table of cluster nodes, you should see your local node, for example `node1.example.com`.

The following is a list of additional information about node IP addresses:

- For the local node only, Oracle Universal Installer (OUI) automatically fills in public and VIP fields. If your system uses vendor clusterware, then OUI may fill additional fields.
- Host names and virtual host names are not domain-qualified. If you provide a domain in the address field during installation, then OUI removes the domain from the address.

- Interfaces identified as private for private IP addresses should not be accessible as public interfaces. Using public interfaces for Cache Fusion can cause performance problems.
- When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `/bin/hostname` command.

a. Click **Add** to add another node to the cluster.

b. Enter the second node's public name (`node2`), and virtual IP name (`node2-vip`), then click **OK**.

You are returned to the Cluster Node Information window. You should now see all nodes listed in the table of cluster nodes. Make sure the Role column is set to HUB for both nodes. To add Leaf Nodes, you must configure GNS.

c. Make sure all nodes are selected, then click the **SSH Connectivity** button at the bottom of the window.

The bottom panel of the window displays the SSH Connectivity information.

d. Enter the operating system user name and password for the Oracle software owner (`grid`). Select the option **If you have configured SSH connectivity between the nodes**, then select the **Reuse private and public keys existing in user home** option. Click **Setup**.

A message window appears, indicating that it might take several minutes to configure SSH connectivity between the nodes. After a short period, another message window appears indicating that passwordless SSH connectivity has been established between the cluster nodes. Click **OK** to continue.

e. When returned to the Cluster Node Information window, click **Next** to continue.

The Specify Network Interface Usage page appears.

8. Select the usage type for each network interface displayed, then click **Next**.

Verify that each interface has the correct interface type associated with it. If you have network interfaces that should not be used by Oracle Clusterware, then set the network interface type to **Do Not Use**. For example, if you have only two network interfaces, then set the public interface to have a Use For value of **Public** and set the private network interface to have a Use For value of **ASM & Private**.

Click **Next**. The ASM Client Storage window appears.

9. Choose the disk group to store Oracle Cluster Registry (OCR) and voting files for the cluster on the Oracle Domain Services Cluster.

Click **Next**.

The Operating System Groups window appears.

10. Accept the default operating system group names for Oracle ASM administration and click **Next**.

The Specify Install Location window appears.

11. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation, then click **Next**. The Oracle base directory must be different from the Oracle home directory.

If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as directed in Step 1, then the default location for the Oracle base directory should display as `/u01/app/grid`.

If you have not installed Oracle software previously on this computer, then the Create Inventory window appears.

12. Change the path for the inventory directory, if required. Then, click **Next**.

If you are using the same directory names as the examples in this book, then it should show a value of `/u01/app/oraInventory`. The group name for the `oraInventory` directory should show `oinstall`.

The Root Script Execution Configuration window appears.

13. Select the option to **Automatically run configuration scripts**. Enter the credentials for the root user or a sudo account, then click **Next**.

Alternatively, you can Run the scripts manually as the `root` user at the end of the installation process when prompted by the installer.

The Perform Prerequisite Checks window appears.

14. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the **Check Again** button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click **Next**.

The Summary window appears.

15. Review the contents of the Summary window and then click **Install**.

The installer displays a progress indicator enabling you to monitor the installation process.

16. If you did not configure automation of the root scripts, then you are required to run certain scripts as the `root` user, as specified in the Execute Configuration Scripts window appears. Do *not* click **OK** until you have run the scripts. Run the scripts on all nodes as directed, in the order shown.

For example, on Oracle Linux you perform the following steps (note that for clarity, the examples show the current user, node and directory in the prompt):

- a. As the `oracle` user on `node1`, open a terminal window, and enter the following commands:

```
[oracle@node1 oracle]$ cd /u01/app/oraInventory
[oracle@node1 oraInventory]$ su
```

- b. Enter the password for the `root` user, and then enter the following command to run the first script on `node1`:

```
[root@node1 oraInventory]# ./oraInstRoot.sh
```

- c. After the `oraInstRoot.sh` script finishes on `node1`, open another terminal window, and as the `oracle` user, enter the following commands:

```
[oracle@node1 oracle]$ ssh node2
[oracle@node2 oracle]$ cd /u01/app/oraInventory
[oracle@node2 oraInventory]$ su
```

- d. Enter the password for the `root` user, and then enter the following command to run the first script on `node2`:

```
[root@node2 oraInventory]# ./oraInstRoot.sh
```

- e. After the `oraInstRoot.sh` script finishes on `node2`, go to the terminal window you opened in part a of this step. As the `root` user on `node1`, enter the following commands to run the second script, `root.sh`:

```
[root@node1 oraInventory]# cd /u01/app/12.2.0/grid
[root@node1 grid]# ./root.sh
```

Press **Enter** at the prompt to accept the default value.

Note:

You must run the `root.sh` script on the first node and wait for it to finish. If your cluster has three or more nodes, then `root.sh` can be run concurrently on all nodes but the first. Node numbers are assigned according to the order of running `root.sh`. If you want to create a particular node number assignment, then run the `root` scripts in the order of the node assignments you want to make, and wait for the script to finish running on each node before proceeding to run the script on the next node. However, Oracle system identifier, or SID, for your Oracle RAC databases, do not follow the node numbers.

- f. After the `root.sh` script finishes on `node1`, go to the terminal window you opened in part c of this step. As the `root` user on `node2`, enter the following commands:

```
[root@node2 oraInventory]# cd /u01/app/12.2.0/grid
[root@node2 grid]# ./root.sh
```

After the `root.sh` script completes, return to the OUI window where the Installer prompted you to run the `oraInstRoot.sh` and `root.sh` scripts. Click **OK**.

The software installation monitoring window reappears.

When you run `root.sh` during Oracle Grid Infrastructure installation, the Trace File Analyzer (TFA) Collector is also installed in the directory `grid_home/tfa`.

17. After `root.sh` runs on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
18. During installation of Oracle Member Cluster for Oracle Databases, if the Member Cluster Manifest file does not include configuration details for Oracle ASM, then Oracle Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM for storage.
19. Continue monitoring the installation until the Finish window appears. Then click **Close** to complete the installation process and exit the installer.

Caution:

After installation is complete, do not remove manually or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running on the server. If you remove these files, then the Oracle software can encounter intermittent hangs. Oracle Clusterware installations can fail with the error:

```
CRS-0184: Cannot communicate with the CRS daemon.
```

After your Oracle Grid Infrastructure installation is complete, you can install Oracle Database on a cluster node for high availability, other applications, or install Oracle RAC.

See Also: *Oracle Real Application Clusters Installation Guide* or *Oracle Database Installation Guide* for your platform for information on installing Oracle Database

Related Topics:

[Creating Member Cluster Manifest File for Oracle Member Clusters](#) (page 8-21)
Create a Member Cluster Manifest file to specify the Oracle Member Cluster configuration for the Grid Infrastructure Management Repository (GIMR), Grid Naming Service, Oracle ASM storage server, and Rapid Home Provisioning configuration.

Related Topics:

[Configuring Device Persistence Manually for Oracle ASM](#) (page B-15)
Complete these tasks to create device path persistence manually for Oracle ASM.

9.4 Installing Oracle Grid Infrastructure Using a Cluster Configuration File

During installation of Oracle Grid Infrastructure, you have the option of either of providing cluster configuration information manually, or of using a cluster configuration file.

A cluster configuration file is a text file that you can create before starting `gridSetup.sh`, which provides the installer with cluster node addresses that it requires to configure the cluster.

Oracle recommends that you consider using a cluster configuration file if you intend to perform repeated installations on a test cluster, or if you intend to perform an installation on many nodes. A sample cluster configuration file is available in the directory `Grid_home/install/response/sample.ccf`.

To create a cluster configuration file manually, start a text editor, and create a file that provides the name of the public and virtual IP addresses for each cluster member node, in the following format:

```
node1 node1-vip /node-role
node2 node2-vip /node-role
.
.
.
```

`node-role` can have either HUB or LEAF as values. Specify the different nodes, separating them with either spaces or colon (:).

For example:

```
mynode1 mynode1-vip /HUB
mynode2 mynode2-vip /LEAF
```

Or, for example:

```
mynode1:mynode1-vip:/HUB
mynode2:mynode2-vip:/LEAF
```

Example 9-1 Sample Cluster Configuration File

The following sample cluster configuration file is available in the directory `Grid_home/install/response/sample.ccf`:

```
#
# Cluster nodes configuration specification file
#
# Format:
# node [vip] [role-identifier] [site-name]
#
# node           - Node's public host name
# vip            - Node's virtual host name
# role-identifier - Node's role with "/" prefix - should be "/HUB" or "/LEAF"
# site-name      - Node's assigned site
#
# Specify details of one node per line.
# Lines starting with '#' will be skipped.
#
# (1) vip and role are not required for Oracle Grid Infrastructure software only
#     installs and Oracle Member cluster for Applications
# (2) vip should be specified as AUTO if Node Virtual host names are Dynamically
#     assigned
# (3) role-identifier can be specified as "/LEAF" only for "Oracle Standalone
#     Cluster"
# (4) site-name should be specified only when configuring Oracle Grid Infrastructure
#     with "Extended Cluster" option
#
# Examples:
# -----
# For installing GI software only on a cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
# node1
# node2
#
# For Standalone Cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
# node1 node1-vip /HUB
# node2 node2-vip /LEAF
#
# For Standalone Extended Cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
# node1 node1-vip /HUB sitea
# node2 node2-vip /LEAF siteb
#
# For Domain Services Cluster:
# ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
# node1 node1-vip /HUB
# node2 node2-vip /HUB
```


selection. You can install Oracle Grid Infrastructure software on multiple nodes at a time.

4. When the software is configured, run the `oraInstRoot.sh` script on all nodes, when prompted.
5. On all nodes, the `root.sh` script output provides information about how to proceed, depending on the configuration you plan to complete in this installation. Make note of this information.

6. Ensure that you have completed all storage and server preinstallation requirements.

7. Verify that all of the cluster nodes meet the installation requirements:

```
runcluvfy.sh stage -pre crsinst -n node_list
```

8. Configure the cluster using the Oracle Universal Installer (OUI) configuration wizard or response files.

Related Topics:

[Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster](#) (page 9-29)

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

[Configuring the Software Binaries Using a Response File](#) (page 9-30)

When you install or copy Oracle Grid Infrastructure software on any node, you can defer configuration for a later time. Review this procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard (`gridSetup.sh`).

[Installing Oracle Domain Services Cluster](#) (page 9-14)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.

[Installing Oracle Standalone Cluster](#) (page 9-7)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.

[Installing Oracle Member Clusters](#) (page 9-21)

Complete this procedure to install Oracle Grid Infrastructure software for Oracle Member Cluster for Oracle Database and Oracle Member Cluster for Applications.

9.5.2 Configuring Software Binaries for Oracle Grid Infrastructure for a Cluster

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

1. Log in on a cluster node as the Oracle Grid Infrastructure installation owner, and change directory to `Grid_home`.
2. Start the Oracle Grid Infrastructure configuration wizard:

```
$ ./gridSetup.sh
```

3. Provide information as needed for configuration. OUI validates the information and configures the installation on all cluster nodes.

4. When you complete providing information, OUI shows you the Summary page, listing the information you have provided for the cluster. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

5. When prompted, run `root` scripts.
6. When you confirm that all `root` scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

9.5.3 Configuring the Software Binaries Using a Response File

When you install or copy Oracle Grid Infrastructure software on any node, you can defer configuration for a later time. Review this procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard (`gridSetup.sh`).

To configure the Oracle Grid Infrastructure software binaries using a response file:

1. As the Oracle Grid Infrastructure installation owner (`grid`), start Oracle Universal Installer in Oracle Grid Infrastructure configuration wizard mode from the Oracle Grid Infrastructure software-only home using the following syntax, where *Grid_home* is the Oracle Grid Infrastructure home, and *filename* is the response file name:

```
Grid_home/gridSetup.sh [-debug] [-silent -responseFile filename]
```

For example:

```
$ cd /u01/app/grid/  
$ ./gridSetup.sh -responseFile /u01/app/grid/response/response_file.rsp
```

The configuration script starts Oracle Universal Installer in Configuration Wizard mode. Each page shows the same user interface and performs the same validation checks that OUI normally does. However, instead of running an installation, the configuration wizard mode validates inputs and configures the installation on all cluster nodes.

2. When you complete configuring values, OUI shows you the Summary page, listing all information you have provided for the cluster. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

3. When prompted, run `root` scripts.
4. When you confirm that all `root` scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

9.5.4 Setting Ping Targets for Network Checks

Receive notification about network status by setting the `Ping_Targets` parameter during the Oracle Grid Infrastructure installation.

For environments where the network link status is not correctly returned when the network cable is disconnected, for example, in a virtual machine, you can receive notification about network status by setting the `Ping_Targets` parameter during the Oracle Grid Infrastructure installation.

Run the installer:

```
./gridSetup.sh oracle_install_crs_Ping_Targets=Host1/IP1,Host2/IP2
```

The ping utility contacts the comma-separated list of host names or IP addresses `Host1/IP1,Host2/IP2` to determine whether the public network is available. If none of the hosts respond, then the network is considered to be offline. Addresses outside the cluster, like of a switch or router, should be used.

For example:

```
/gridSetup.sh oracle_install_crs_Ping_Targets=192.0.2.1,192.0.2.2
```

9.6 About Deploying Oracle Grid Infrastructure Using Rapid Home Provisioning

Rapid Home Provisioning is a software lifecycle management method for provisioning and patching Oracle homes. Rapid Home Provisioning enables mass deployment of standard operating environments for databases and clusters.

Rapid Home Provisioning (RHP) enables you to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes. The supported versions are 11.2, 12.1, and 12.2. You can also provision applications and middleware using Rapid Home Provisioning. A single cluster, known as the Rapid Home Provisioning Server, stores and manages standardized images, called gold images, which can be provisioned to any number of nodes. You can install Oracle Grid Infrastructure cluster configurations such as Oracle Standalone Clusters, Oracle Member Clusters, and Oracle Application Clusters. After deployment, you can expand and contract clusters and Oracle RAC Databases.

You can provision Oracle Grid Infrastructure on a remote set of nodes in a cloud computing environment from a single cluster where you store templates of Oracle homes as images (called gold images) of Oracle software, such as databases, middleware, and applications.

Rapid Home Provisioning leverages a new file system capability which allows for separation of gold image software from the site-specific configuration changes, so the home path remains unchanged throughout updates. This capability of persistent home path is available for Oracle Grid Infrastructure 12c Release 2 (12.2) and combines the benefits of in-place and out-of-place patching.

Note: Rapid Home Provisioning supports provisioning, patching, and upgrade of single-instance databases on Oracle Grid Infrastructure for a standalone server, or Oracle Restart.

Rapid Home Provisioning

Deploying Oracle software using Rapid Home Provisioning has the following advantages:

- Ensures standardization and enables high degrees of automation with gold images and managed lineage of deployed software.

- Supports change management. With standardized Oracle homes, an administrator has better control of the hosted Oracle software and can easily manage the mass deployment and maintenance of the software through a single location for change management.
- Minimizes downtime during patching and upgrades, eases rollbacks, and makes provisioning for large systems easier and more efficient.
- Reduces the cumulative time to patch software images, since a single Oracle home may be used for many database instances.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about setting up the Rapid Home Provisioning Server and Client, creating and using gold images for provisioning and patching Oracle Grid Infrastructure and Oracle Database homes.

9.7 Confirming Oracle Clusterware Function

After Oracle Grid Infrastructure installation, confirm that your Oracle Clusterware installation is installed and running correctly.

After installation, log in as `root`, and use the following command syntax to confirm that your Oracle Clusterware installation is installed and running correctly:

```
crsctl check cluster -all
```

For example:

```
$ crsctl check cluster -all

*****
node1:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
*****
node2:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
*****
node3:
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
*****
```

Note: After installation is complete, do not remove manually or run cron jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` or its files while Oracle Clusterware is up. If you remove these files, then Oracle Clusterware could encounter intermittent hangs, and you will encounter error `CRS-0184: Cannot communicate with the CRS daemon`.

9.8 Confirming Oracle ASM Function for Oracle Clusterware Files

Confirm Oracle ASM is running after installing Oracle Grid Infrastructure.

After Oracle Grid Infrastructure installation, Oracle Clusterware files are stored on Oracle ASM. Use the following command syntax as the Oracle Grid Infrastructure installation owner (`grid`) to confirm that your Oracle ASM installation is running:

```
srvctl status asm
```

For example:

```
srvctl status asm  
ASM is running on node1,node2, node3, node4
```

Note: To manage Oracle ASM or Oracle Net 11g Release 2 (11.2) or later installations, use the `srvctl` binary in the Oracle Grid Infrastructure home for a cluster (Grid home). If you have Oracle Real Application Clusters or Oracle Database installed, then you cannot use the `srvctl` binary in the database home to manage Oracle ASM or Oracle Net.

9.9 Understanding Offline Processes in Oracle Grid Infrastructure

After the installation of Oracle Grid Infrastructure, some components may be listed as OFFLINE. Oracle Grid Infrastructure activates these resources when you choose to add them.

Oracle Grid Infrastructure provides required resources for various Oracle products and components. Some of those products and components are optional, so you can install and enable them after installing Oracle Grid Infrastructure. To simplify postinstall additions, Oracle Grid Infrastructure preconfigures and registers all required resources for all products available for these products and components, but only activates them when you choose to add them. As a result, some components may be listed as OFFLINE after the installation of Oracle Grid Infrastructure.

Resources listed as TARGET:OFFLINE and STATE:OFFLINE do not need to be monitored. They represent components that are registered, but not enabled, so they do not use any system resources. If an Oracle product or component is installed on the system, and it requires a particular resource to be online, then the software prompts you to activate the required offline resource.

Oracle Grid Infrastructure Postinstallation Tasks

Complete configuration task after you install Oracle Grid Infrastructure.

You are required to complete some configuration tasks after Oracle Grid Infrastructure is installed. In addition, Oracle recommends that you complete additional tasks immediately after installation. You must also complete product-specific configuration tasks before you use those products.

Note: This chapter describes basic configuration only. Refer to product-specific administration and tuning guides for more detailed configuration and tuning information.

[Required Postinstallation Tasks](#) (page 10-1)

Download and apply required patches for your software release after completing your initial installation.

[Recommended Postinstallation Tasks](#) (page 10-2)

Oracle recommends that you complete these tasks after installation.

[About Changes in Default SGA Permissions for Oracle Database](#) (page 10-7)

Starting with Oracle Database 12c Release 2 (12.2.0.1), by default, permissions to read and write to the System Global Area (SGA) are limited to the Oracle software installation owner.

[Using Earlier Oracle Database Releases with Oracle Grid Infrastructure](#) (page 10-7)

Review the following topics for information about using earlier Oracle Database releases with Oracle Grid Infrastructure 12c Release 2 (12.2) installations:

[Modifying Oracle Clusterware Binaries After Installation](#) (page 10-10)

After installation, if you need to modify the Oracle Clusterware configuration, then you must unlock the Grid home. Review this information about unlocking the Grid home.

10.1 Required Postinstallation Tasks

Download and apply required patches for your software release after completing your initial installation.

[Downloading and Installing Patch Updates](#) (page 10-2)

Download and install patch updates for your Oracle software after you complete installation.

10.1.1 Downloading and Installing Patch Updates

Download and install patch updates for your Oracle software after you complete installation.

Check the My Oracle Support website for required patch updates for your installation.

1. Use a web browser to view the My Oracle Support website:

<https://support.oracle.com>

2. Log in to My Oracle Support website.

Note: If you are not a My Oracle Support registered user, then click **Register for My Oracle Support** and register.

3. On the main My Oracle Support page, click **Patches & Updates**.
4. In the **Patch Search** region, select **Product or Family (Advanced)**.
5. On the **Product or Family (Advanced)** display, provide information about the product, release, and platform for which you want to obtain patches, and click **Search**.

The Patch Search pane opens, displaying the results of your search.

6. Select the patch number and click **ReadMe**.

The README page is displayed. It contains information about the patch set and how to apply the patches to your installation.

7. Use the unzip utility provided with the software to uncompress the Oracle patch updates that you downloaded from My Oracle Support. The unzip utility is located in the `$ORACLE_HOME/bin` directory.

10.2 Recommended Postinstallation Tasks

Oracle recommends that you complete these tasks after installation.

[Tuning Semaphore Parameters](#) (page 10-3)

Refer to the following guidelines if the default semaphore parameter values are too low to accommodate all Oracle processes.

[Creating a Backup of the root.sh Script](#) (page 10-3)

Oracle recommends that you back up the `root.sh` script after you complete an installation.

[Downloading and Installing the ORAchk Health Check Tool](#) (page 10-3)

Download and install the ORAchk utility to perform proactive health checks for the Oracle software stack.

[Creating a Fast Recovery Area](#) (page 10-4)

During an Oracle Restart installation, you can create only one disk group. During an Oracle Clusterware installation, you can create multiple disk groups. If you plan to add an Oracle Database for a standalone server or an Oracle RAC database, then you should create the fast recovery area for database files.

[Checking the SCAN Configuration](#) (page 10-5)

The Single Client Access Name (SCAN) is a name that is used to provide service access for clients to the cluster. Because the SCAN is associated with the cluster as a whole, rather than to a particular node, the SCAN makes it possible to add or remove nodes from the cluster without needing to reconfigure clients.

[Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications](#) (page 10-6)

After you have completed Oracle Grid Infrastructure installation, you can set resource limits in the `Grid_home/crs/install/s_crsconfig_nodename_env.txt` file.

10.2.1 Tuning Semaphore Parameters

Refer to the following guidelines if the default semaphore parameter values are too low to accommodate all Oracle processes.

Note:

Oracle recommends that you refer to the operating system documentation for more information about setting semaphore parameters.

1. Calculate the minimum total semaphore requirements using the following formula:

$$2 * \text{sum}(\text{process parameters of all database instances on the system}) + \text{overhead for background processes} + \text{system and other application requirements}$$
2. Set `semms` (total semaphores systemwide) to this total.
3. Set `semms1` (semaphores for each set) to 250.
4. Set `semmsi` (total semaphore sets) to `semms` divided by `semms1`, rounded up to the nearest multiple of 1024.

10.2.2 Creating a Backup of the root.sh Script

Oracle recommends that you back up the `root.sh` script after you complete an installation.

If you install other products in the same Oracle home directory subsequent to this installation, then Oracle Universal Installer updates the contents of the existing `root.sh` script during the installation. If you require information contained in the original `root.sh` script, then you can recover it from the backed up `root.sh` file.

10.2.3 Downloading and Installing the ORAchk Health Check Tool

Download and install the ORAchk utility to perform proactive health checks for the Oracle software stack.

ORAchk replaces the RACCheck utility. ORAchk extends health check coverage to the entire Oracle software stack, and identifies and addresses top issues reported by Oracle users. ORAchk proactively scans for known problems with Oracle products and deployments, including the following:

- Standalone Oracle Database
- Oracle Grid Infrastructure
- Oracle Real Application Clusters
- Maximum Availability Architecture (MAA) Validation
- Upgrade Readiness Validations
- Oracle Golden Gate

Oracle is continuing to expand checks, based on customer requests.

ORAchk is supported on Windows Server 2012 and Windows Server 2016 on a Cygwin environment only.

Oracle recommends that you download and run the latest version of ORAchk from My Oracle Support. For information about downloading, configuring and running ORAchk utility, refer to My Oracle Support note 1268927.2:

<https://support.oracle.com/epmos/faces/DocContentDisplay?id=1268927.2&parent=DOCUMENTATION&sourceId=USERGUIDE>

Related Topics:

Oracle ORAchk and EXAchk User's Guide

10.2.4 Creating a Fast Recovery Area

During an Oracle Restart installation, you can create only one disk group. During an Oracle Clusterware installation, you can create multiple disk groups. If you plan to add an Oracle Database for a standalone server or an Oracle RAC database, then you should create the fast recovery area for database files.

[About the Fast Recovery Area and the Fast Recovery Area Disk Group](#)
(page 10-4)

The fast recovery area is a unified storage location for all Oracle Database files related to recovery. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

[Creating the Fast Recovery Area Disk Group](#) (page 10-5)

Procedure to create the fast recovery area disk group.

10.2.4.1 About the Fast Recovery Area and the Fast Recovery Area Disk Group

The fast recovery area is a unified storage location for all Oracle Database files related to recovery. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

Database administrators can define the `DB_RECOVERY_FILE_DEST` parameter to the path for the fast recovery area to enable on disk backups and rapid recovery of data. When you enable fast recovery in the `init.ora` file, Oracle Database writes all RMAN backups, archive logs, control file automatic backups, and database copies to the fast recovery area. RMAN automatically manages files in the fast recovery area by deleting obsolete backups and archiving files no longer required for recovery.

Oracle recommends that you create a fast recovery area disk group. Oracle Clusterware files and Oracle Database files can be placed on the same disk group, and you can also place fast recovery files in the same disk group. However, Oracle

recommends that you create a separate fast recovery disk group to reduce storage device contention.

The fast recovery area is enabled by setting the `DB_RECOVERY_FILE_DEST` parameter. The size of the fast recovery area is set with `DB_RECOVERY_FILE_DEST_SIZE`. As a general rule, the larger the fast recovery area, the more useful it becomes. For ease of use, Oracle recommends that you create a fast recovery area disk group on storage devices that can contain at least three days of recovery information. Ideally, the fast recovery area is large enough to hold a copy of all of your data files and control files, the online redo logs, and the archived redo log files needed to recover your database using the data file backups kept under your retention policy.

Multiple databases can use the same fast recovery area. For example, assume you have created a fast recovery area disk group on disks with 150 GB of storage, shared by 3 different databases. You can set the size of the fast recovery for each database depending on the importance of each database. For example, if database1 is your least important database, database2 is of greater importance, and database3 is of greatest importance, then you can set different `DB_RECOVERY_FILE_DEST_SIZE` settings for each database to meet your retention target for each database: 30 GB for database1, 50 GB for database2, and 70 GB for database3.

10.2.4.2 Creating the Fast Recovery Area Disk Group

Procedure to create the fast recovery area disk group.

1. Go to the `Grid_home/bin` directory, and start Oracle ASM Configuration Assistant (ASMCA).

For example:

```
$ cd /u01/app/oracle/product/12.2.0/grid/bin
$ ./asmca
```

ASMCA opens at the Disk Groups tab.

2. Click **Create** to create a new disk group.

The Create Disk Groups window opens.

3. Provide configuration information for the fast recovery area as prompted:

In the Disk Group Name field, enter a descriptive name for the fast recovery area group. For example: `FRA`.

In the Redundancy section, select the level of redundancy you want to use. For example: `Normal`.

In the Select Member Disks field, select eligible disks you want to add to the fast recovery area, and click **OK**.

The Diskgroup Creation window opens and provides disk group creation status.

4. When the Fast Recovery Area disk group creation is complete, click **OK**, and then click **Exit**.

10.2.5 Checking the SCAN Configuration

The Single Client Access Name (SCAN) is a name that is used to provide service access for clients to the cluster. Because the SCAN is associated with the cluster as a whole,

rather than to a particular node, the SCAN makes it possible to add or remove nodes from the cluster without needing to reconfigure clients.

The Single Client Access Name (SCAN) also adds location independence for the databases, so that client configuration does not have to depend on which nodes are running a particular database instance. Clients can continue to access the cluster in the same way as with previous releases, but Oracle recommends that clients accessing the cluster use the SCAN.

You can use the command `cluvfy comp scan` (located in *Grid home/bin*) to confirm that the DNS is correctly associating the SCAN with the addresses. For example:

```
$cluvfy comp scan

Verifying Single Client Access Name (SCAN) ...
  Verifying DNS/NIS name service 'rws127064-clu-scan.rws127064-
clu.rws12706410644.example.com' ...
    Verifying Name Service Switch Configuration File Integrity ...PASSED
  Verifying DNS/NIS name service 'rws127064-clu-scan.rws127064-
clu.rws12706410644.example.com' ...PASSED
Verifying Single Client Access Name (SCAN) ...PASSED

Verification of SCAN was successful.

CVU operation performed: SCAN
Date: Jul 29, 2016 1:42:41 AM
CVU home: /u01/crshome/
User: crsusr
```

After installation, when a client sends a request to the cluster, the Oracle Clusterware SCAN listeners redirect client requests to servers in the cluster.

See Also:

Oracle Clusterware Administration and Deployment Guide for more information about system checks and configurations

10.2.6 Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications

After you have completed Oracle Grid Infrastructure installation, you can set resource limits in the *Grid_home/crs/install/s_crsconfig_nodename_env.txt* file.

The resource limits apply to all Oracle Clusterware processes and Oracle databases managed by Oracle Clusterware. For example, to set a higher number of processes limit, edit the file and set the `CRS_LIMIT_NPROC` parameter to a high value.

```
---
#Do not modify this file except as documented above or under the
#direction of Oracle Support Services.
#####
TZ=PST8PDT
NLS_LANG=AMERICAN_AMERICA.WE8ISO8859P1
CRS_LIMIT_STACK=2048
CRS_LIMIT_OPENFILE=65536
CRS_LIMIT_NPROC=65536
TNS_ADMIN=
```

10.3 About Changes in Default SGA Permissions for Oracle Database

Starting with Oracle Database 12c Release 2 (12.2.0.1), by default, permissions to read and write to the System Global Area (SGA) are limited to the Oracle software installation owner.

In previous releases, both the Oracle installation owner account and members of the OSDBA group had access to shared memory. The change in Oracle Database 12c Release 2 (12.2) to restrict access by default to the Oracle installation owner account provides greater security than previous configurations. However, this change may prevent DBAs who do not have access to the Oracle installation owner account from administering the database.

The Oracle Database initialization parameter `ALLOW_GROUP_ACCESS_TO_SGA` determines if the Oracle Database installation owner account (`oracle` in Oracle documentation examples) is the only user that can read and write to the database System Global Area (SGA), or if members of the OSDBA group can read the SGA. In Oracle Database 12c Release 2 (12.2), the default value for this parameter is `FALSE`, so that only the Oracle Database installation owner has read and write permissions to the SGA. Group access to the SGA is removed by default. This change affects all Linux and UNIX platforms.

If members of the OSDBA group require read access to the SGA, then you can change the initialization parameter `ALLOW_GROUP_ACCESS_TO_SGA` setting from `FALSE` to `TRUE`. Oracle strongly recommends that you accept the default permissions that limit access to the SGA to the `oracle` user account.

Related Topics:

Oracle Database Reference

10.4 Using Earlier Oracle Database Releases with Oracle Grid Infrastructure

Review the following topics for information about using earlier Oracle Database releases with Oracle Grid Infrastructure 12c Release 2 (12.2) installations:

[General Restrictions for Using Earlier Oracle Database Releases](#) (page 10-8)

You can use Oracle Database 12c releases 1 and 2 and Oracle Database 11g release 2 (11.2.0.3 or later) with Oracle Grid Infrastructure 12c release 2 (12.2).

[Configuring Earlier Release Oracle Database on Oracle ACFS](#) (page 10-8)

Review this information to configure a 11.2 release Oracle Database on Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

[Managing Server Pools with Earlier Database Versions](#) (page 10-9)

Starting with Oracle Grid Infrastructure 12c, Oracle Database server categories include roles such as Hub and Leaf that were not present in earlier releases.

[Making Oracle ASM Available to Earlier Oracle Database Releases](#) (page 10-10)

To use Oracle ASM with Oracle Database releases earlier than Oracle Database 12c, you must use Local ASM or set the cardinality for Oracle Flex ASM to `ALL`, instead of the default of 3.

[Using ASMCA to Administer Disk Groups for Earlier Database Releases](#) (page 10-10)

Use Oracle ASM Configuration Assistant (ASMCA) to create and modify disk groups when you install earlier Oracle databases and Oracle RAC databases on Oracle Grid Infrastructure installations.

[Using the Correct LSNRCTL Commands](#) (page 10-10)

To administer Oracle Database 12c Release 2 local and scan listeners using the `lsnrctl` command, set your `$ORACLE_HOME` environment variable to the path for the Oracle Grid Infrastructure home (Grid home).

10.4.1 General Restrictions for Using Earlier Oracle Database Releases

You can use Oracle Database 12c releases 1 and 2 and Oracle Database 11g release 2 (11.2.0.3 or later) with Oracle Grid Infrastructure 12c release 2 (12.2).

Do not use the versions of `srvctl`, `lsnrctl`, or other Oracle Grid infrastructure home tools to administer earlier version databases. Only administer earlier Oracle Database releases using the tools in the earlier Oracle Database homes. To ensure that the versions of the tools you are using are the correct tools for those earlier release databases, run the tools from the Oracle home of the database or object you are managing.

Oracle Database homes can only be stored on Oracle ASM Cluster File System (Oracle ACFS) if the database version is Oracle Database 11g release 2 or later. Earlier releases of Oracle Database cannot be installed on Oracle ACFS because these releases were not designed to use Oracle ACFS.

When installing 11.2 databases on an Oracle Flex ASM cluster, the Oracle ASM cardinality must be set to `All`.

Note:

If you are installing Oracle Database 11g release 2 with Oracle Grid Infrastructure 12c release 2 (12.2), then before running Oracle Universal Installer (OUI) for Oracle Database, run the following command on the local node only:

```
Grid_home/oui/bin/runInstaller -ignoreSysPrereqs -updateNodeList  
ORACLE_HOME=Grid_home "CLUSTER_NODES={comma_separated_list_of_hub_nodes}"  
CRS=true LOCAL_NODE=local_node [-cfs]
```

Use the `-cfs` option only if the `Grid_home` is on a shared location.

10.4.2 Configuring Earlier Release Oracle Database on Oracle ACFS

Review this information to configure a 11.2 release Oracle Database on Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

1. Install Oracle Grid Infrastructure 12c Release 2 (12.2) as described in this guide.
2. Start Oracle ASM Configuration Assistant (ASMCA) as the grid installation owner. For example:

```
./asmca
```

Follow the steps in the configuration wizard to create Oracle ACFS storage for the earlier release Oracle Database home.

3. Install Oracle Database 11g release 2 (11.2) software-only on the Oracle ACFS file system you configured.
4. From the 11.2 Oracle Database home, run Oracle Database Configuration Assistant (DBCA) and create the Oracle RAC Database, using Oracle ASM as storage for the database data files.

```
./dbca
```

5. Modify the Oracle ACFS path dependency:

```
srvctl modify database -d my_112_db -j Oracle_ACFS_path
```

Related Topics:

Oracle Automatic Storage Management Administrator's Guide

10.4.3 Managing Server Pools with Earlier Database Versions

Starting with Oracle Grid Infrastructure 12c, Oracle Database server categories include roles such as Hub and Leaf that were not present in earlier releases.

For this reason, you cannot create server pools using the Oracle RAC 11g Release 2 (11.2) version of Database Configuration Assistant (DBCA). To create server pools for earlier release Oracle RAC installations, use the following procedure.

1. Log in as the Oracle Grid Infrastructure installation owner (Grid user).
2. Change directory to the 12.2 Oracle Grid Infrastructure binaries directory in the Grid home. For example:

```
# cd /u01/app/12.2.0/grid/bin
```

3. Use the Oracle Grid Infrastructure 12c version of `srvctl` to create a server pool consisting of Hub Node roles. For example, to create a server pool called `p_hub` with a maximum size of one cluster node, enter the following command:

```
srvctl add serverpool -serverpool p_hub -min 0 -max 1 -category hub;
```

4. Log in as the Oracle RAC installation owner, start DBCA from the Oracle RAC Oracle home. For example:

```
$ cd /u01/app/oracle/product/12.2.0/dbhome_1/bin
$ dbca
```

DBCA discovers the server pool that you created with the Oracle Grid Infrastructure 12c `srvctl` command. Configure the server pool as required for your services.

See Also:

Oracle Clusterware Administration and Deployment Guide for more information about managing resources using policies

10.4.4 Making Oracle ASM Available to Earlier Oracle Database Releases

To use Oracle ASM with Oracle Database releases earlier than Oracle Database 12c, you must use Local ASM or set the cardinality for Oracle Flex ASM to ALL, instead of the default of 3.

After you install Oracle Grid Infrastructure 12c, if you want to use Oracle ASM to provide storage service for Oracle Database releases that are earlier than Oracle Database 12c, then you must use the following command to modify the Oracle ASM resource (`ora.asm`):

```
$ srvctl modify asm -count ALL
```

This setting changes the cardinality of the Oracle ASM resource so that Oracle Flex ASM instances run on all cluster nodes. You must change the setting even if you have a cluster with three or less than three nodes to ensure database releases earlier than 11g Release 2 can find the `ora.node.sid.inst` resource alias.

10.4.5 Using ASMCA to Administer Disk Groups for Earlier Database Releases

Use Oracle ASM Configuration Assistant (ASMCA) to create and modify disk groups when you install earlier Oracle databases and Oracle RAC databases on Oracle Grid Infrastructure installations.

Starting with Oracle Database 11g Release 2, Oracle ASM is installed as part of an Oracle Grid Infrastructure installation, with Oracle Clusterware. You can no longer use Database Configuration Assistant (DBCA) to perform administrative tasks on Oracle ASM.

See Also:

Oracle Automatic Storage Management Administrator's Guide for details about configuring disk group compatibility for databases using Oracle Database 11g or earlier software with Oracle Grid Infrastructure 12c (12.2)

10.4.6 Using the Correct LSNRCTL Commands

To administer Oracle Database 12c Release 2 local and scan listeners using the `lsnrctl` command, set your `$ORACLE_HOME` environment variable to the path for the Oracle Grid Infrastructure home (Grid home).

Do not attempt to use the `lsnrctl` commands from Oracle home locations for previous releases, as they cannot be used with the new release.

10.5 Modifying Oracle Clusterware Binaries After Installation

After installation, if you need to modify the Oracle Clusterware configuration, then you must unlock the Grid home. Review this information about unlocking the Grid home.

For example, if you want to apply a one-off patch, or if you want to modify an Oracle Exadata configuration to run IPC traffic over RDS on the interconnect instead of using the default UDP, then you must unlock the Grid home.

Caution:

Before relinking executables, you must shut down all executables that run in the Oracle home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries.

Unlock the home using the following procedure:

1. Change directory to the path *Grid_home/crs/install*, where *Grid_home* is the path to the Grid home, and unlock the Grid home using the command `rootcrs.sh -unlock`. For example, with the Grid home `/u01/app/12.2.0/grid`, enter the following command:

```
# cd /u01/app/12.2.0/grid/crs/install
# rootcrs.sh -unlock
```

2. Change user to the Oracle Grid Infrastructure software owner, and relink binaries using the command syntax `make -f Grid_home/rdbms/lib/ins_rdbms.mk target`, where *Grid_home* is the Grid home, and *target* is the binaries that you want to relink. For example, where the grid user is `grid`, `$ORACLE_HOME` is set to the Grid home, and where you are updating the interconnect protocol from UDP to IPC, enter the following command:

```
# su grid
$ make -f $ORACLE_HOME/rdbms/lib/ins_rdbms.mk ipc_rds ioracle
```

Note:

To relink binaries, you can also change to the grid installation owner and run the command `Grid_home/bin/relink`.

3. Relock the Grid home and restart the cluster as follows:

```
# rootcrs.sh -lock
# crsctl start crs
```

Repeat steps 1 through 3 on each cluster member node.

Note:

Do not delete directories in the Grid home. For example, do not delete the directory `Grid_home/OPatch`. If you delete the directory, then the Grid infrastructure installation owner cannot use OPatch to patch the Grid home, and OPatch displays the error message "checkdir error: cannot create `Grid_home/OPatch`".

Upgrading Oracle Grid Infrastructure

Oracle Grid Infrastructure upgrade consists of upgrade of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

Oracle Grid Infrastructure upgrades can be rolling upgrades, in which a subset of nodes are brought down and upgraded while other nodes remain active. Starting with Oracle ASM 12c release 2 (12.2), upgrades can be rolling upgrades.

You can also use Rapid Home Provisioning to upgrade Oracle Grid Infrastructure for a cluster.

[Understanding Out-of-Place Upgrade](#) (page 11-2)

Review this information about out-of-place upgrade of Oracle Grid Infrastructure.

[About Oracle Grid Infrastructure Upgrade and Downgrade](#) (page 11-3)

Review this information about upgrade and downgrade of Oracle Grid Infrastructure.

[Options for Oracle Grid Infrastructure Upgrades](#) (page 11-3)

Understand the upgrade options for Oracle Grid Infrastructure in this release. When you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2), you upgrade to an Oracle Flex Cluster configuration.

[Restrictions for Oracle Grid Infrastructure Upgrades](#) (page 11-4)

Review the following information for restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

[Preparing to Upgrade an Existing Oracle Clusterware Installation](#) (page 11-6)

If you have an existing Oracle Clusterware installation, then you upgrade your existing cluster by performing an out-of-place upgrade. You cannot perform an in-place upgrade.

[Understanding Rolling Upgrades Using Batches](#) (page 11-12)

Review this information to understand rolling upgrade of Oracle Grid Infrastructure.

[Performing Rolling Upgrade of Oracle Grid Infrastructure](#) (page 11-13)

Review this information to perform rolling upgrade of Oracle Grid Infrastructure.

[About Upgrading Oracle Grid Infrastructure Using Rapid Home Provisioning](#) (page 11-17)

Rapid Home Provisioning is a software lifecycle management method for provisioning and patching Oracle homes.

[Applying Patches to Oracle Grid Infrastructure](#) (page 11-18)

After you have upgraded Oracle Grid Infrastructure 12c Release 2 (12.2), you can install individual software patches by downloading them from My Oracle Support.

[Updating Oracle Enterprise Manager Cloud Control Target Parameters](#) (page 11-19)

After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.

[Unlocking the Existing Oracle Clusterware Installation](#) (page 11-21)

After upgrade from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure Grid home, then you must first change the permission and ownership of the previous release Grid home.

[Checking Cluster Health Monitor Repository Size After Upgrading](#) (page 11-22)

If you are upgrading Oracle Grid Infrastructure from a prior release using IPD/OS to the current release, then review the Cluster Health Monitor repository size (the CHM repository).

[Downgrading Oracle Clusterware After an Upgrade](#) (page 11-23)

After a successful or a failed upgrade, you can restore Oracle Clusterware to the previous release.

[Completing Failed or Interrupted Installations and Upgrades](#) (page 11-29)

If Oracle Universal Installer (OUI) exits on the node from which you started the upgrade, or the node reboots before you confirm that the `rootupgrade.sh` script was run on all nodes, then the upgrade remains incomplete.

[Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure](#) (page 11-32)

Review this information to convert to an Oracle Extended Cluster after upgrading Oracle Grid Infrastructure. Oracle Extended Cluster enables you to deploy Oracle RAC databases on a cluster, in which some of the nodes are located in different sites.

Related Topics:

[About Upgrading Oracle Grid Infrastructure Using Rapid Home Provisioning](#) (page 11-17)

Rapid Home Provisioning is a software lifecycle management method for provisioning and patching Oracle homes.

11.1 Understanding Out-of-Place Upgrade

Review this information about out-of-place upgrade of Oracle Grid Infrastructure.

With an out-of-place upgrade, the installer installs the newer version in a separate Oracle Clusterware home. Both versions of Oracle Clusterware are on each cluster member node, but only one version is active.

Rolling upgrade avoids downtime and ensure continuous availability while the software is upgraded to a new version.

If you have separate Oracle Clusterware homes on each node, then you can perform an out-of-place upgrade on all nodes, or perform an out-of-place rolling upgrade, so that some nodes are running Oracle Clusterware from the earlier version Oracle

Clusterware home, and other nodes are running Oracle Clusterware from the new Oracle Clusterware home.

An in-place upgrade of Oracle Grid Infrastructure is not supported.

11.2 About Oracle Grid Infrastructure Upgrade and Downgrade

Review this information about upgrade and downgrade of Oracle Grid Infrastructure.

You can upgrade Oracle Grid Infrastructure in any of the following ways:

- Rolling Upgrade which involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster
- Non-rolling Upgrade which involves bringing down all the nodes except one. A complete cluster outage occurs while the root script stops the old Oracle Clusterware stack and starts the new Oracle Clusterware stack on the node where you initiate the upgrade. After upgrade is completed, the new Oracle Clusterware is started on all the nodes.

Note that some services are disabled when one or more nodes are in the process of being upgraded. All upgrades are out-of-place upgrades, meaning that the software binaries are placed in a different Grid home from the Grid home used for the prior release.

You can downgrade from Oracle Grid Infrastructure 12c Release 2 (12.2) to Oracle Grid Infrastructure 12c Release 1 (12.1) and Oracle Grid Infrastructure 11g Release 2 (11.2). Be aware that if you downgrade to a prior release, then your cluster must conform with the configuration requirements for that prior release, and the features available for the cluster consist only of the features available for that prior release of Oracle Clusterware and Oracle ASM.

You can perform out-of-place upgrades to an Oracle ASM instance using Oracle ASM Configuration Assistant (ASMCA). In addition to running ASMCA using the graphical user interface, you can run ASMCA in non-interactive (silent) mode.

Note:

You must complete an upgrade before attempting to use cluster backup files. You cannot use backups for a cluster that has not completed upgrade.

See Also:

Oracle Database Upgrade Guide and *Oracle Automatic Storage Management Administrator's Guide* for additional information about upgrading existing Oracle ASM installations

11.3 Options for Oracle Grid Infrastructure Upgrades

Understand the upgrade options for Oracle Grid Infrastructure in this release. When you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2), you upgrade to an Oracle Flex Cluster configuration.

Supported upgrade paths for Oracle Grid Infrastructure for this release are:

- Oracle Grid Infrastructure upgrade from releases 11.2.0.3 and 11.2.0.4 to Oracle Grid Infrastructure 12c Release 2 (12.2).

- Oracle Grid Infrastructure upgrade from Oracle Grid Infrastructure 12c Release 1 (12.1) to Oracle Grid Infrastructure 12c Release 2 (12.2).

Upgrade options from Oracle Grid Infrastructure 11g and Oracle Grid Infrastructure 12c Release 1 (12.1) to Oracle Grid Infrastructure 12c Release 2 (12.2) include the following:

- Oracle Grid Infrastructure rolling upgrade which involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster
- Oracle Grid Infrastructure non-rolling upgrade by bringing the cluster down and upgrading the complete cluster

Note:

- When you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2), you upgrade to an Oracle Standalone Cluster configuration.
 - If storage for OCR and voting files is other than Oracle ASM, you need to migrate OCR and voting files to Oracle ASM before upgrading to Oracle Grid Infrastructure 12c Release 2 (12.2).
-
-

11.4 Restrictions for Oracle Grid Infrastructure Upgrades

Review the following information for restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

- Oracle Grid Infrastructure upgrades are always out-of-place upgrades. You cannot perform an in-place upgrade of Oracle Grid Infrastructure to existing homes.
- The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 12c Release 2 (12.2) upgrade.
- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.
- When you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2), you upgrade to an Oracle Flex Cluster configuration.
- Do not delete directories in the Grid home. For example, do not delete the directory *Grid_home*/Opatch. If you delete the directory, then the Grid infrastructure installation owner cannot use Opatch to patch the grid home, and Opatch displays the error message "'checkdir' error: cannot create *Grid_home*/OPatch".
- To upgrade existing Oracle Grid Infrastructure installations to Oracle Grid Infrastructure 12c Release 2 (12.2), you must first verify if you need to apply any mandatory patches for upgrade to succeed.

Oracle recommends that you use the Cluster Verification Utility tool (CVU) to check if there are any patches required for upgrading your existing Oracle Grid Infrastructure or Oracle RAC database installations. See *Using CVU to Validate Readiness for Oracle Clusterware Upgrades* for steps to check readiness.

- The software in the 12c Release 2 (12.2) Oracle Grid Infrastructure home is not fully functional until the upgrade is completed. Running `srvctl`, `crsctl`, and other commands from the new Grid homes are not supported until the final `rootupgrade.sh` script is run and the upgrade is complete across all nodes.

To manage databases in existing earlier release database homes during the Oracle Grid Infrastructure upgrade, use the `srvctl` from the existing database homes.

- To change a cluster member node role to Leaf, you must have completed the upgrade on all Oracle Grid Infrastructure nodes so that the active version is Oracle Grid Infrastructure 12c Release 1 (12.1) or later.
- To upgrade existing Oracle Clusterware installations to Oracle Grid Infrastructure 12c cluster, your release must be greater than or equal to Oracle Grid Infrastructure 11g Release 2 (11.2.0.3).

See Also:

Oracle Database Upgrade Guide for additional information about preparing for upgrades

About Storage Restrictions for Upgrade

- If the Oracle Cluster Registry (OCR) and voting file locations for your current installation are on raw or block devices, or shared file systems, then you must migrate them to Oracle ASM disk groups before upgrading to Oracle Grid Infrastructure 12c Release 2 (12.2).
- If you want to upgrade Oracle Grid Infrastructure releases before Oracle Grid Infrastructure 11g Release 2 (11.2), where the OCR and voting files are on raw or block devices or shared file system, then you must upgrade to Oracle Grid Infrastructure 11g Release 2 (11.2), and move the Oracle Cluster Registry (OCR) and voting files to Oracle ASM, before you upgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).
- If you have Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems on Oracle Grid Infrastructure 11g Release 2 (11.2.0.1), you upgrade Oracle Grid Infrastructure to any later release, and you take advantage of Redundant Interconnect Usage and add one or more additional private interfaces to the private network, then you must restart the Oracle ASM instance on each upgraded cluster member node.

About Upgrading Shared Grid Homes

- If the existing Oracle Clusterware home is a shared home, then you can use a non-shared home for the Oracle Grid Infrastructure for a cluster home for Oracle Clusterware and Oracle ASM 12c Release 2 (12.2).
- You can perform upgrades on a shared Oracle Clusterware home.

About Single-Instance Oracle ASM Upgrade

- During Oracle Grid Infrastructure installation or upgrade, if there is a single instance Oracle ASM release on the local node, then it is converted to an Oracle

Flex ASM 12c Release 2 (12.2) installation, and Oracle ASM runs in the Oracle Grid Infrastructure home on all nodes.

- If a single instance (non-clustered) Oracle ASM installation is on a remote node, which is a node other than the local node (the node on which the Oracle Grid Infrastructure installation or upgrade is being performed), then it remains a single instance Oracle ASM installation. However, during the installation or upgrade, when the OCR and voting files are placed on Oracle ASM, then an Oracle Flex ASM installation is created on all nodes in the cluster. The single instance Oracle ASM installation on the remote node becomes nonfunctional.

Related Topics:

[Using CVU to Validate Readiness for Oracle Clusterware Upgrades](#)

(page 11-11)

Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

11.5 Preparing to Upgrade an Existing Oracle Clusterware Installation

If you have an existing Oracle Clusterware installation, then you upgrade your existing cluster by performing an out-of-place upgrade. You cannot perform an in-place upgrade.

The following topics list the steps you can perform before you upgrade Oracle Grid Infrastructure:

[Upgrade Checklist for Oracle Grid Infrastructure](#) (page 11-6)

Review this checklist before upgrading an existing Oracle Grid Infrastructure. A cluster is being upgraded until all cluster member nodes are running the new installations, and the new clusterware becomes the active version.

[Checks to Complete Before Upgrading Oracle Grid Infrastructure](#) (page 11-8)

Complete the following tasks before upgrading Oracle Grid Infrastructure.

[Moving Oracle Clusterware Files from NFS to Oracle ASM](#) (page 11-9)

If Oracle Cluster Registry (OCR) and voting files are stored on Network File System (NFS), then move these files to Oracle ASM disk groups before upgrading Oracle Grid Infrastructure.

[Running the Oracle ORAchK Upgrade Readiness Assessment](#) (page 11-10)

Download and run the ORAchK Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.

[Using CVU to Validate Readiness for Oracle Clusterware Upgrades](#)

(page 11-11)

Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

11.5.1 Upgrade Checklist for Oracle Grid Infrastructure

Review this checklist before upgrading an existing Oracle Grid Infrastructure. A cluster is being upgraded until all cluster member nodes are running the new installations, and the new clusterware becomes the active version.

Table 11-1 Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
Review Upgrade Guide for deprecation and desupport information that may affect upgrade planning.	<i>Oracle Database Upgrade Guide</i>
Patch set (recommended)	Install the latest patch set release for your existing installation. Review My Oracle Support note 2180188.1 for the list of latest patches before upgrading Oracle Grid Infrastructure.
Install user account	Confirm that the installation owner you plan to use is the same as the installation owner that owns the installation you want to upgrade.
Create a Grid home	Create a new Oracle Grid Infrastructure Oracle home (Grid home) where you can extract the image files. All Oracle Grid Infrastructure upgrades (upgrades of existing Oracle Clusterware and Oracle ASM installations) are out-of-place upgrades.
Instance names for Oracle ASM	Oracle Automatic Storage Management (Oracle ASM) instances must use standard Oracle ASM instance names. The default ASM SID for a single-instance database is +ASM.
Cluster names and Site names	Cluster names must have the following characteristics: <ul style="list-style-type: none"> • At least one character but no more than 15 characters in length. • Hyphens (-), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9). • It cannot begin with a numeric character. • It cannot begin or end with the hyphen (-) character.
Operating System	Confirm that you are using a supported operating system, kernel release, and all required operating system packages for the new Oracle Grid Infrastructure installation.
Network addresses for standard Oracle Grid Infrastructure	For standard Oracle Grid Infrastructure installations, confirm the following network configuration: <ul style="list-style-type: none"> • The private and public IP addresses are in unrelated, separate subnets. The private subnet should be in a dedicated private subnet. • The public and virtual IP addresses, including the SCAN addresses, are in the same subnet (the range of addresses permitted by the subnet mask for the subnet network). • Neither private nor public IP addresses use a link local subnet (169.254.*.*).
OCR on raw or block devices	Migrate OCR files from RAW or Block devices to Oracle ASM or a supported file system. Direct use of RAW and Block devices is not supported. Run the <code>ocrcheck</code> command to confirm Oracle Cluster Registry (OCR) file integrity. If this check fails, then repair the OCR before proceeding.
CVU Upgrade Validation	Use Cluster Verification Utility (CVU) to assist you with system checks in preparation for starting an upgrade.
Unset Environment variables	As the user performing the upgrade, unset the environment variables <code>\$ORACLE_HOME</code> and <code>\$ORACLE_SID</code> . Check that the <code>ORA_CRS_HOME</code> environment variable is not set. Do not use <code>ORA_CRS_HOME</code> as an environment variable, except under explicit direction from Oracle Support. Refer to the <i>Checks to Complete Before Upgrading Oracle Grid Infrastructure</i> for a complete list of environment variables to unset.

Table 11-1 (Cont.) Upgrade Checklist for Oracle Grid Infrastructure Installation

Check	Task
RACcheck Upgrade Readiness Assessment	Download and run the RACcheck Upgrade Readiness Assessment to obtain automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure. See My Oracle Support note 1457357.1, which is available at the following URL: https://support.oracle.com/rs?type=doc&id=1457357.1
Back Up the Oracle Software Before Upgrades	Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle software and databases.
HugePages memory allocation	Allocate memory to HugePages large enough for the System Global Areas (SGA) of all databases planned to run on the cluster, and to accommodate the System Global Area for the Grid Infrastructure Management Repository.

Related Topics:

[Checks to Complete Before Upgrading Oracle Grid Infrastructure](#) (page 11-8)
Complete the following tasks before upgrading Oracle Grid Infrastructure.

Related Topics:

[Moving Oracle Clusterware Files from NFS to Oracle ASM](#) (page 11-9)
If Oracle Cluster Registry (OCR) and voting files are stored on Network File System (NFS), then move these files to Oracle ASM disk groups before upgrading Oracle Grid Infrastructure.

[Checks to Complete Before Upgrading Oracle Grid Infrastructure](#) (page 11-8)
Complete the following tasks before upgrading Oracle Grid Infrastructure.

[My Oracle Support Note 2180188.1](#)

11.5.2 Checks to Complete Before Upgrading Oracle Grid Infrastructure

Complete the following tasks before upgrading Oracle Grid Infrastructure.

1. For each node, use Cluster Verification Utility to ensure that you have completed preinstallation steps. It can generate Fixup scripts to help you to prepare servers. In addition, the installer helps you to ensure all required prerequisites are met.

Ensure that you have the information you need during installation, including the following:

- An Oracle base location for Oracle Clusterware.
- An Oracle Grid Infrastructure home location that is different from your existing Oracle Clusterware location.
- SCAN name and addresses, and other network addresses.
- Privileged user operating system groups.
- `root` user access, to run scripts as `root` during installation.

- For the installation owner running the installation, if you have environment variables set for the existing installation, then unset the environment variables `$ORACLE_HOME` and `$ORACLE_SID`, as these environment variables are used during upgrade. For example, as `grid` user, run the following commands on the local node:

For bash shell:

```
$ unset ORACLE_BASE
$ unset ORACLE_HOME
$ unset ORACLE_SID
```

For C shell:

```
$ unsetenv ORACLE_BASE
$ unsetenv ORACLE_HOME
$ unsetenv ORACLE_SID
```

- If you have set `ORA_CRS_HOME` as an environment variable, following instructions from Oracle Support, then unset it before starting an installation or upgrade. You should never use `ORA_CRS_HOME` as an environment variable except under explicit direction from Oracle Support.
- Check to ensure that the user profile for the installation user, for example, `.profile` or `.cshrc`, does not set any of these environment variables.
- If you have an existing installation on your system, and you are using the same user account to install this installation, then unset the following environment variables: `ORA_CRS_HOME`, `ORACLE_HOME`, `ORA_NLS10`, `TNS_ADMIN` and any other environment variable set for the Oracle installation user that is connected with Oracle software homes.
- Ensure that the `$ORACLE_HOME/bin` path is removed from your `PATH` environment variable.

Related Topics:

[Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle Database](#) (page 6-1)

Before installation, create operating system groups and users, and configure user environments.

[Configuring Networks for Oracle Grid Infrastructure and Oracle RAC](#) (page 5-1)

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

11.5.3 Moving Oracle Clusterware Files from NFS to Oracle ASM

If Oracle Cluster Registry (OCR) and voting files are stored on Network File System (NFS), then move these files to Oracle ASM disk groups before upgrading Oracle Grid Infrastructure.

- As Oracle Grid Infrastructure installation owner (`grid`), create the Oracle ASM disk group using `ASMCA`.

```
./asmca
```

Follow the steps in the ASMCA wizard to create the Oracle ASM disk group, for example, DATA.

2. As grid user, move the voting files to the Oracle ASM disk group you created:

```
crsctl replace votedisk +DATA
```

The output of this command is as follows:

```
CRS-4256: Updating the profile
Successful addition of voting disk 24c6d682874a4f1ebf54f5ab0098b9e4.
Successful deletion of voting disk 1b5044fa39684f86bfbe681f388e55fb.
Successfully replaced voting disk group with +DATA_DG_OCR_VDSK.
CRS-4256: Updating the profile
CRS-4266: Voting file(s) successfully replaced
```

3. As grid user, check the Oracle Cluster Registry (OCR) status:

```
./ocrcheck
```

The output of the command is as follows:

```
Status of Oracle Cluster Registry is as follows :
Version           :          4
Total space (kbytes) :    409568
Used space (kbytes)  :       1380
Available space (kbytes) :    408188
ID                :   288871063
Device/File Name   : /oradbcfs/storage/12101/ocr
Device/File integrity check succeeded
Cluster registry integrity check succeeded
```

4. As root user, move the OCR files to the Oracle ASM disk group you created:

```
./ocrconfig -add +DATA
```

5. As root user, delete the Oracle Clusterware files from the NFS location:

```
./ocrconfig -delete ocr_file_path_previously_on_nfs
```

11.5.4 Running the Oracle ORAchk Upgrade Readiness Assessment

Download and run the ORAchk Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.

ORAchk is an Oracle RAC configuration audit tool. ORAchk Upgrade Readiness Assessment can be used to obtain an automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure 11.2.0.3, 11.2.0.4, 12.1.0.1, 12.1.0.2, and 12.2. You can run the ORAchk Upgrade Readiness Assessment tool and automate many of the manual pre-upgrade and post-upgrade checks.

Oracle recommends that you download and run the latest version of ORAchk from My Oracle Support. For information about downloading, configuring, and running ORAchk, refer to My Oracle Support note 1457357.1.

See Also:

- <https://support.oracle.com/rs?type=doc&id=1457357.1>
 - *Oracle ORAchk and EXAchk User's Guide*
-
-

11.5.5 Using CVU to Validate Readiness for Oracle Clusterware Upgrades

Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

You can use CVU to assist you with system checks in preparation for starting an upgrade. CVU runs the appropriate system checks automatically, and either prompts you to fix problems, or provides a fixup script to be run on all nodes in the cluster before proceeding with the upgrade.

[About the CVU Upgrade Validation Command Options](#) (page 11-11)

Review this information about running upgrade validations.

[Example of Verifying System Upgrade Readiness for Grid Infrastructure](#) (page 11-12)

You can verify that the permissions required for installing Oracle Clusterware have been configured on the nodes `node1` and `node2` by running a command similar to the following.

11.5.5.1 About the CVU Upgrade Validation Command Options

Review this information about running upgrade validations.

- Run Oracle Universal Installer (OUI), and allow the Cluster Verification Utility (CVU) validation built into OUI to perform system checks and generate fixup scripts.
- Run the CVU manual script `cluvfy.sh` to perform system checks and generate fixup scripts.

To use OUI to perform pre-install checks and generate fixup scripts, run the installation as you normally would. OUI starts CVU, and performs system checks as part of the installation process. Selecting OUI to perform these checks is particularly appropriate if you think you have completed preinstallation checks, and you want to confirm that your system configuration meets minimum requirements for installation.

To use the `cluvfy.sh` command-line script for CVU, navigate to the new Grid home where you extracted the image files for upgrade, that contains the `runcluvfy.sh` script, and run the command `runcluvfy.sh stage -pre crsinst -upgrade` to check the readiness of your Oracle Clusterware installation for upgrades. Running `runcluvfy.sh` with the `-pre crsinst -upgrade` options performs system checks to confirm if the cluster is in a correct state for upgrading from an existing clusterware installation.

The command uses the following syntax, where variable content is indicated by italics:

```
runcluvfy.sh stage -pre crsinst -upgrade [-rolling]
-src_crshome src_Gridhome ]-dest_crshome dest_Gridhome -dest_version dest_release
[-fixup][-fixupnoexec][-method sudo -user user_name [-location dir_path][-method
root][-verbose]
```

The options are:

- `-rolling`
Use this option to verify readiness for rolling upgrades.
- `-src_crshome src_Gridhome`

Use this option to indicate the location of the source Oracle Clusterware or Grid home that you are upgrading, where *src_Gridhome* is the path to the home that you want to upgrade.

- `-dest_crshome dest_Gridhome`

Use this option to indicate the location of the upgrade Grid home, where *dest_Gridhome* is the path to the Grid home.

- `-dest_version dest_release`

Use the `-dest_version` option to indicate the release number of the upgrade, including any patchset. The release number must include the five digits designating the release to the level of the platform-specific patch. For example: 12.2.0.1.0.

- `-fixup [-method sudo -user user_name [-location dir_path] [-method root]`

Use the `-fixup` option to indicate that you want to generate instructions for any required steps you need to complete to ensure that your cluster is ready for an upgrade. The default location is the CVU work directory.

The `-fixup -method` option defines the method by which root scripts are run. The `-method` flag requires one of the following options:

- `sudo`: Run as a user on the sudoers list.
- `root`: Run as the root user.

If you select `sudo`, then enter the `-location` option to provide the path to Sudo on the server, and enter the `-user` option to provide the user account with Sudo privileges.

- `-fixupnoexec`

If the option is specified, then on verification failure, the fix up data is generated and the instruction for manual execution of the generated fix ups is displayed.

- `-verbose`

Use the `-verbose` flag to produce detailed output of individual checks.

11.5.5.2 Example of Verifying System Upgrade Readiness for Grid Infrastructure

You can verify that the permissions required for installing Oracle Clusterware have been configured on the nodes `node1` and `node2` by running a command similar to the following.

```
$ ./runcluvfy.sh stage -pre crsinst -upgrade -rolling -src_crshome
/u01/app/11.2.0/grid -dest_crshome /u01/app/12.2.0/grid -dest_version
12.2.0.1 -fixup -verbose
```

See Also:

Oracle Database Upgrade Guide

11.6 Understanding Rolling Upgrades Using Batches

Review this information to understand rolling upgrade of Oracle Grid Infrastructure.

When you upgrade Oracle Grid Infrastructure, you upgrade the entire cluster. You cannot select or de-select individual nodes for upgrade. Oracle does not support attempting to add additional nodes to a cluster during a rolling upgrade. Oracle recommends that you leave Oracle RAC instances running when upgrading Oracle Clusterware. When you start the `root` script on each node, the database instances on that node are shut down and then the `rootupgrade.sh` script starts the instances again.

You can use `root` user automation to automate running the `rootupgrade.sh` script during the upgrade. When you use `root` automation, you can divide the nodes into groups, or batches, and start upgrades of these batches. Between batches, you can move services from nodes running the previous release to the upgraded nodes, so that services are not affected by the upgrade. Oracle recommends that you use `root` automation, and allow the `rootupgrade.sh` script to stop and start instances automatically. You can also continue to run `root` scripts manually.

Restrictions for Selecting Nodes for Batch Upgrades

The following restrictions apply when selecting nodes in batches for upgrade:

- You can pool nodes in batches for upgrade, up to a maximum of three batches.
- The local node, where Oracle Universal Installer (OUI) is running, must be upgraded in batch one.
- Hub and Leaf Nodes cannot be upgraded in the same batch.
- All Hub Nodes must be upgraded before starting the upgrade of Leaf Nodes.

11.7 Performing Rolling Upgrade of Oracle Grid Infrastructure

Review this information to perform rolling upgrade of Oracle Grid Infrastructure.

[Upgrading Oracle Grid Infrastructure from an Earlier Release](#) (page 11-13)

Complete this procedure to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) from an earlier release.

[Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable](#) (page 11-16)

If some nodes become unreachable in the middle of an upgrade, then you cannot complete the upgrade, because the upgrade script (`rootupgrade.sh`) did not run on the unreachable nodes. Because the upgrade is incomplete, Oracle Clusterware remains in the previous release.

[Joining Inaccessible Nodes After Forcing an Upgrade](#) (page 11-16)

Use this procedure to join inaccessible nodes after a force cluster upgrade.

[Changing the First Node for Install and Upgrade](#) (page 11-16)

If the first node becomes inaccessible, you can force another node to be the first node for installation or upgrade.

11.7.1 Upgrading Oracle Grid Infrastructure from an Earlier Release

Complete this procedure to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) from an earlier release.

At any time during the upgrade, if you have a question about what you are being asked to do, or what input you are required to provide during upgrade, click the **Help** button on the installer page.

You should have your network information, storage information, and operating system users and groups available to you before you start upgrade, and you should be prepared to run root scripts.

1. As `grid` user, download the Oracle Grid Infrastructure image files and extract the files to the Grid home.

For example:

```
mkdir -p /u01/app/12.2.0/grid
chown grid:oinstall /u01/app/12.2.0/grid
cd /u01/app/12.2.0/grid
unzip -q download_location/grid_home.zip
```

`download_location/grid_home.zip` is the path of the downloaded Oracle Grid Infrastructure image file.

Note:

- You must extract the image software into the directory where you want your Grid home to be located.
 - Download and copy the Oracle Grid Infrastructure image files to the local node only. During upgrade, the software is copied and installed on all other nodes in the cluster.
-
-

2. Start the Oracle Grid Infrastructure wizard by running the following command:

```
Grid_home/gridSetup.sh
```

3. Select the following configuration option:

- **Upgrade Oracle Grid Infrastructure:** Select this option to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).
-
-

Note: Oracle Clusterware must always be the later release, so you cannot upgrade Oracle ASM to a release that is more recent than Oracle Clusterware.

4. On the Node Selection page, select all nodes.
5. Select installation options as prompted. Oracle recommends that you configure `root` script automation, so that the `rootupgrade.sh` script can be run automatically during the upgrade.
6. Run `root` scripts, using either automatically or manually:
 - **Running root scripts automatically:**

If you have configured `root` script automation, then use the pause between batches to relocate services from the nodes running the previous release to the new release.
 - **Running root scripts manually**

If you have not configured `root` script automation, then when prompted, run the `rootupgrade.sh` script on each node in the cluster that you want to upgrade.

If you run `root` scripts manually, then run the script on the local node first. The script shuts down the earlier release installation, replaces it with the new Oracle Clusterware release, and starts the new Oracle Clusterware installation. After the script completes successfully, you can run the script in parallel on all nodes except for one, which you select as the last node. When the script is run successfully on all the nodes except the last node, run the script on the last node. When upgrading from 12.1 Oracle Flex Cluster, Oracle recommends that you run the `rootupgrade.sh` script on all Hub Nodes before running it on Leaf Nodes.

7. Because the Oracle Grid Infrastructure home is in a different location than the former Oracle Clusterware and Oracle ASM homes, update any scripts or applications that use utilities, libraries, or other files that reside in the Oracle Clusterware and Oracle ASM homes.
8. Update the Oracle Enterprise Manager target parameters as described in the topic *Updating Oracle Enterprise Manager Cloud Control Target Parameters*.

Note:

- At the end of the upgrade, if you set the Oracle Cluster Registry (OCR) backup location manually to the earlier release Oracle Clusterware home (CRS home), then you must change the OCR backup location to the new Oracle Grid Infrastructure home (Grid home). If you did not set the OCR backup location manually, then the backup location is changed for you during the upgrade.
- Because upgrades of Oracle Clusterware are out-of-place upgrades, the previous release Oracle Clusterware home cannot be the location of the current release OCR backups. Backups in the old Oracle Clusterware home can be deleted.
- If the cluster being upgraded has a single disk group that stores the OCR, OCR backup, Oracle ASM password, Oracle ASM password file backup, and the Grid Infrastructure Management Repository (GIMR), then Oracle recommends that you create a separate disk group or use another existing disk group and store the OCR backup, the GIMR and Oracle ASM password file backup in that disk group.

See Also: *Oracle Clusterware Administration and Deployment Guide* for the commands to create a disk group.

Related Topics:

[Updating the Enterprise Manager Cloud Control Target After Upgrades](#)
(page 11-20)

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

11.7.2 Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable

If some nodes become unreachable in the middle of an upgrade, then you cannot complete the upgrade, because the upgrade script (`rootupgrade.sh`) did not run on the unreachable nodes. Because the upgrade is incomplete, Oracle Clusterware remains in the previous release.

You can confirm that the upgrade is incomplete by entering the command `crsctl query crs activeversion`.

To resolve this problem, run the `rootupgrade.sh` command with the `-force` flag using the following syntax:

```
Grid_home/rootupgrade -force
```

For example:

```
# /u01/app/12.2.0/grid/rootupgrade -force
```

This command forces the upgrade to complete. Verify that the upgrade has completed by using the command `crsctl query crs activeversion`. The active release should be the upgrade release.

The force cluster upgrade has the following limitations:

- All active nodes must be upgraded to the newer release
- All inactive nodes (accessible or inaccessible) may be either upgraded or not upgraded
- For inaccessible nodes, after patch set upgrades, you can delete the node from the cluster. If the node becomes accessible later, and the patch version upgrade path is supported, then you can upgrade it to the new patch version.

11.7.3 Joining Inaccessible Nodes After Forcing an Upgrade

Use this procedure to join inaccessible nodes after a force cluster upgrade.

Starting with Oracle Grid Infrastructure 12c, after you complete a force cluster upgrade, you can use the procedure described here to join inaccessible nodes to the cluster as an alternative to deleting the nodes, which was required in earlier releases. To use this option, you must already have Oracle Grid Infrastructure 12c Release 2 (12.2) software installed on the nodes.

1. Log in as the `root` user on the node that you want to join to the cluster.
2. Change directory to the Oracle Grid Infrastructure 12c release 2 (12.2) `Grid_home` directory. For example:

```
$ cd /u01/12.2.0/grid/
```

3. Run the following command, where `upgraded_node` is the inaccessible or unreachable node that you want to join to the cluster:

```
$ rootupgrade.sh -join -existingnode upgraded_node
```

11.7.4 Changing the First Node for Install and Upgrade

If the first node becomes inaccessible, you can force another node to be the first node for installation or upgrade.

During installation, if `root.sh` fails to complete on the first node, run the following command on another node using the `-force` option:

```
root.sh -force -first
```

For upgrade:

```
rootupgrade.sh -force -first
```

11.8 About Upgrading Oracle Grid Infrastructure Using Rapid Home Provisioning

Rapid Home Provisioning is a software lifecycle management method for provisioning and patching Oracle homes.

Rapid Home Provisioning (RHP) enables you to install clusters, and provision, patch, and upgrade Oracle Grid Infrastructure and Oracle Database homes. The supported versions are 11.2, 12.1, and 12.2. You can also provision applications and middleware using Rapid Home Provisioning. A single cluster, known as the Rapid Home Provisioning Server, stores and manages standardized images, called gold images, which can be provisioned to any number of nodes. You can install Oracle Grid Infrastructure cluster configurations such as Oracle Standalone Clusters, Oracle Member Clusters, and Oracle Application Clusters. After deployment, you can expand and contract clusters and Oracle RAC Databases.

You can provision Oracle Grid Infrastructure on a remote set of nodes in a cloud computing environment from a single cluster where you store templates of Oracle homes as images (called gold images) of Oracle software, such as databases, middleware, and applications.

Note: Rapid Home Provisioning is not supported for provisioning, patching, or upgrade of Oracle Grid Infrastructure for a standalone server, or Oracle Restart.

Rapid Home Provisioning

Deploying Oracle software using Rapid Home Provisioning has the following advantages:

- Ensures standardization and enables high degrees of automation with gold images and managed lineage of deployed software.
- Supports change management. With standardized Oracle homes, an administrator has better control of the hosted Oracle software and can easily manage the mass deployment and maintenance of the software through a single location for change management.
- Minimizes downtime during patching and upgrades, eases rollbacks, and makes provisioning for large systems easier and more efficient.
- Reduces the cumulative time to patch software images, since a single Oracle home may be used for many database instances.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about setting up the Rapid Home Provisioning Server and Client, creating and using gold images for provisioning and patching Oracle Grid Infrastructure and Oracle Database homes.

11.9 Applying Patches to Oracle Grid Infrastructure

After you have upgraded Oracle Grid Infrastructure 12c Release 2 (12.2), you can install individual software patches by downloading them from My Oracle Support.

[About Individual \(One-Off\) Oracle Grid Infrastructure Patches](#) (page 11-18)

Download Oracle ASM one-off patch and apply it to Oracle Grid Infrastructure using the OPatch Utility.

[About Oracle Grid Infrastructure Software Patch Levels](#) (page 11-18)

Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.

[Patching Oracle Grid Infrastructure to a Software Patch Level](#) (page 11-19)

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), a new cluster state called "Rolling Patch" is available. This mode is similar to the existing "Rolling Upgrade" mode in terms of the Oracle ASM operations allowed in this quiesce state.

11.9.1 About Individual (One-Off) Oracle Grid Infrastructure Patches

Download Oracle ASM one-off patch and apply it to Oracle Grid Infrastructure using the OPatch Utility.

Individual patches are called **one-off** patches. An Oracle ASM one-off patch is available for a specific released release of Oracle ASM. If a patch you want is available, then you can download the patch and apply it to Oracle ASM using the OPatch Utility. The OPatch inventory keeps track of the patches you have installed for your release of Oracle ASM. If there is a conflict between the patches you have installed and patches you want to apply, then the OPatch Utility advises you of these conflicts.

11.9.2 About Oracle Grid Infrastructure Software Patch Levels

Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.

The software patch level for Oracle Grid Infrastructure represents the set of all one-off patches applied to the Oracle Grid Infrastructure software release, including Oracle ASM. The release is the release number, in the format of major, minor, and patch set release number. For example, with the release number 12.1.0.1, the major release is 12, the minor release is 1, and 0.0 is the patch set number. With one-off patches, the major and minor release remains the same, though the patch levels change each time you apply or roll back an interim patch.

As with standard upgrades to Oracle Grid Infrastructure, at any given point in time for normal operation of the cluster, all the nodes in the cluster must have the same software release and patch level. Because one-off patches can be applied as rolling upgrades, all possible patch levels on a particular software release are compatible with each other.

11.9.3 Patching Oracle Grid Infrastructure to a Software Patch Level

Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), a new cluster state called "Rolling Patch" is available. This mode is similar to the existing "Rolling Upgrade" mode in terms of the Oracle ASM operations allowed in this quiesce state.

1. Download patches you want to apply from My Oracle Support:

<https://support.oracle.com>

Select the Patches and Updates tab to locate the patch.

Oracle recommends that you select **Recommended Patch Advisor**, and enter the product group, release, and platform for your software. My Oracle Support provides you with a list of the most recent patch set updates (PSUs) and critical patch updates (CPUs).

Place the patches in an accessible directory, such as /tmp.

2. Change directory to the /opatch directory in the Grid home. For example:

```
$ cd /u01/app/12.2.0/grid/opatch
```

3. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.
4. Follow the instructions in the patch documentation to apply the patch.

11.10 Updating Oracle Enterprise Manager Cloud Control Target Parameters

After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.

Because Oracle Grid Infrastructure 12c Release 2 (12.2) is an out-of-place upgrade of the Oracle Clusterware home in a new location (the Oracle Grid Infrastructure for a cluster home, or Grid home), the path for the CRS_HOME parameter in some parameter files must be changed. If you do not change the parameter, then you encounter errors such as "cluster target broken" on Oracle Enterprise Manager Cloud Control.

To resolve the issue, update the Enterprise Manager Cloud Control target, and then update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

[Updating the Enterprise Manager Cloud Control Target After Upgrades](#) (page 11-20)

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

[Updating the Enterprise Manager Agent Base Directory After Upgrades](#) (page 11-20)

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

[Registering Resources with Oracle Enterprise Manager After Upgrades](#) (page 11-20)

After upgrading Oracle Grid Infrastructure, add the new resource targets to Oracle Enterprise Manager Cloud Control.

11.10.1 Updating the Enterprise Manager Cloud Control Target After Upgrades

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

1. Log in to Enterprise Manager Cloud Control.
2. Navigate to the **Targets** menu, and then to the **Cluster** page.
3. Click a cluster target that was upgraded.
4. Click **Cluster**, then **Target Setup**, and then **Monitoring Configuration** from the menu.
5. Update the value for **Oracle Home** with the new Grid home path.
6. Save the updates.

11.10.2 Updating the Enterprise Manager Agent Base Directory After Upgrades

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

The Agent Base directory is a directory where the Management Agent home is created. The Management Agent home is in the path *Agent_Base_Directory/core/EMAgent_Version*. For example, if the Agent Base directory is `/u01/app/emagent`, then the Management Agent home is created as `/u01/app/emagent/core/13.1.1.0`.

1. Navigate to the `bin` directory in the Management Agent home.
2. In the `/u01/app/emagent/core/13.1.1.0/bin` directory, open the file `emctl` with a text editor.
3. Locate the parameter `CRS_HOME`, and update the parameter to the new Grid home path.
4. Repeat steps 1-3 on each node of the cluster with an Enterprise Manager agent.

11.10.3 Registering Resources with Oracle Enterprise Manager After Upgrades

After upgrading Oracle Grid Infrastructure, add the new resource targets to Oracle Enterprise Manager Cloud Control.

Discover and add new resource targets in Oracle Enterprise Manager after Oracle Grid Infrastructure upgrade. The following procedure provides an example of discovering an Oracle ASM listener target after upgrading Oracle Grid Infrastructure.

1. Log in to Oracle Enterprise Manager Cloud Control.
2. From the **Setup** menu, select **Add Target**, and then select **Add Targets Manually**.
The Add Targets Manually page is displayed.
3. In the Add Targets page, select the **Add Using Guided Process** option and **Target Type** as `Oracle Database, Listener and Automatic Storage Management`.

For any other resource to be added, select the appropriate **Target Type** in Oracle Enterprise Manager discovery wizard.

4. Click Add Using Guided Process.

The Target Discover wizard is displayed.

5. For the Specify Host or Cluster field, click on the Search icon and search for Target Types of Hosts, and select the corresponding Host.

6. Click Next.

7. In the Target Discovery: Results page, select the discovered Oracle ASM Listener target, and click Configure.

8. In the Configure Listener dialog box, specify the listener properties and click OK.

9. Click Next and complete the discovery process.

The listener target is discovered in Oracle Enterprise Manager with the status as Down.

10. From the Targets menu, select the type of target.

11. Click the target name to navigate to the target home page.

12. From the host, database, middleware target, or application menu displayed on the target home page, select Target Setup, then select Monitoring Configuration.

13. In the Monitoring Configuration page for the listener, specify the host name in the Machine Name field and the password for the ASMSNMP user in the Password field.

14. Click OK.

Oracle ASM listener target is displayed with the correct status.

Similarly, you can add other clusterware resources to Oracle Enterprise Manager after an Oracle Grid Infrastructure upgrade.

11.11 Unlocking the Existing Oracle Clusterware Installation

After upgrade from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure Grid home, then you must first change the permission and ownership of the previous release Grid home.

Unlock the Oracle Clusterware installation using the following procedure:

1. Log in as `root`, and change the permission and ownership of the previous release Grid home using the following command syntax, where `oldGH` is the previous release Grid home, `swowner` is the Oracle Grid Infrastructure installation owner, and `oldGHParent` is the parent directory of the previous release Grid home:

```
#chmod -R 755 oldGH
#chown -R swowner oldGH
#chown swowner oldGHParent
```

For example:

```
#chmod -R 755 /u01/app/11.2.0/grid
#chown -R grid /u01/app/11.2.0/grid
#chown grid /u01/app/11.2.0
```

2. After you change the permissions and ownership of the previous release Grid home, log in as the Oracle Grid Infrastructure Installation owner (`grid`, in the preceding example), and use the same release Oracle Grid Infrastructure 12c standalone deinstallation tool to remove the previous release Grid home (`oldGH`).

Caution: You must use the deinstallation tool from the same release to remove Oracle software. Do not run the deinstallation tool from a later release to remove Oracle software from an earlier release. For example, do not run the deinstallation tool from the 12.2.0.1 installation media to remove Oracle software from an existing 12.1.0.2 Oracle home.

You can obtain the standalone deinstallation tool from the following URL:

<http://www.oracle.com/technetwork/database/enterprise-edition/downloads/index.html>

Click the **See All** link for the downloads for your operating system platform, and scan the list of downloads for the deinstall utility.

Related Topics:

[About Oracle Deinstallation Options](#) (page 12-2)

You can stop and remove Oracle Database software and components in an Oracle Database home with Oracle Universal Installer.

11.12 Checking Cluster Health Monitor Repository Size After Upgrading

If you are upgrading Oracle Grid Infrastructure from a prior release using IPD/OS to the current release, then review the Cluster Health Monitor repository size (the CHM repository).

1. Review your CHM repository needs, and determine if you need to increase the repository size to maintain a larger CHM repository.

Note: Your previous IPD/OS repository is deleted when you install Oracle Grid Infrastructure.

By default, the CHM repository size is a minimum of either 1GB or 3600 seconds (1 hour), regardless of the size of the cluster.

2. To enlarge the CHM repository, use the following command syntax, where `RETENTION_TIME` is the size of CHM repository in number of seconds:

```
oclumon manage -repos changeretentiontime RETENTION_TIME
```

For example, to set the repository size to four hours:

```
oclumon manage -repos changeretentiontime 14400
```

The value for `RETENTION_TIME` must be more than 3600 (one hour) and less than 259200 (three days). If you enlarge the CHM repository size, then you must ensure that there is local space available for the repository size you select on each node of the cluster. If you do not have sufficient space available, then you can move the repository to shared storage.

11.13 Downgrading Oracle Clusterware After an Upgrade

After a successful or a failed upgrade, you can restore Oracle Clusterware to the previous release.

Downgrading Oracle Clusterware restores the Oracle Clusterware configuration to the state it was in before the Oracle Grid Infrastructure 12c Release 2 (12.2) upgrade. Any configuration changes you performed during or after the Oracle Grid Infrastructure 12c Release 2 (12.2) upgrade are removed and cannot be recovered.

To restore Oracle Clusterware to the previous release, use the downgrade procedure for the release to which you want to downgrade.

Note: Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), you can downgrade the cluster nodes in any sequence. You can downgrade all cluster nodes except one, in parallel. You must downgrade the last node after you downgrade all other nodes.

Note: When downgrading after a failed upgrade, if the `rootcrs.sh` or `rootcrs.bat` file does not exist on a node, then instead of the executing the script use the command `perl rootcrs.pl`. Use the perl interpreter located in the Oracle Home directory.

[Options for Oracle Grid Infrastructure Downgrades](#) (page 11-23)

Understand the downgrade options for Oracle Grid Infrastructure in this release.

[Restrictions for Oracle Grid Infrastructure Downgrades](#) (page 11-24)

Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.

[Downgrading to Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#) (page 11-24)

Use this procedure to downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).

[Downgrading to Oracle Grid Infrastructure 11g Release 2 \(11.2\)](#) (page 11-26)

Use this procedure to downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).

[Downgrading Oracle Grid Infrastructure after Upgrade Fails](#) (page 11-27)

If upgrade of Oracle Grid Infrastructure fails before setting the active version of Oracle Clusterware, then follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

[Downgrading Oracle Grid Infrastructure after Upgrade Fails on Remote Nodes](#) (page 11-28)

If upgrade of Oracle Grid Infrastructure fails on remote nodes, then you can follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

11.13.1 Options for Oracle Grid Infrastructure Downgrades

Understand the downgrade options for Oracle Grid Infrastructure in this release.

Downgrade options from Oracle Grid Infrastructure 12c to earlier releases include the following:

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2). Because all cluster configurations in Oracle Grid Infrastructure 12c Release 2 (12.2) are Oracle Flex Clusters, when you downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2), you downgrade from an Oracle Flex cluster configuration to a Standard cluster configuration.

Related Topics:

[My Oracle Support Note 2180188.1](#)

11.13.2 Restrictions for Oracle Grid Infrastructure Downgrades

Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.

- When you downgrade from an Oracle Grid Infrastructure 12c Release 2 (12.2) to Oracle Grid Infrastructure 11g Release 2 (11.2), you downgrade from an Oracle Flex cluster configuration to a Standard cluster configuration since all cluster configurations in releases earlier than Oracle Grid Infrastructure 12c are Standard cluster configurations. Leaf nodes from the Oracle Grid Infrastructure 12c Release 2 (12.2) cluster are not be a part of the Oracle Grid Infrastructure 11g Release 2 (11.2) standard cluster after the downgrade.
- You can only downgrade to the Oracle Grid Infrastructure release you upgraded from. For example, if you upgraded from Oracle Grid Infrastructure 11g Release 2 (11.2) to Oracle Grid Infrastructure 12c Release 1 (12.2), you can only downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).
- If the cluster has Hub and Leaf Nodes, then the last node to be downgraded must be a Hub Node.

11.13.3 Downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1)

Use this procedure to downgrade to Oracle Grid Infrastructure 12c Release 1 (12.1).

1. Delete the Oracle Grid Infrastructure 12c Release 2 (12.2) Management Database:

```
dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. Use the command syntax `rootcrs.sh -downgrade` to downgrade Oracle Grid Infrastructure on all nodes, in any sequence. For example:

```
# /u01/app/12.2.0/grid/crs/install/rootcrs.sh -downgrade
```

Run this command from a directory that has write permissions for the Oracle Grid Infrastructure installation user. You can run the downgrade script in parallel on all cluster nodes, but one.

3. Downgrade the last node after you downgrade all other nodes:

```
# /u01/app/12.2.0/grid/crs/install/rootcrs.sh -downgrade
```

4. Remove Oracle Grid Infrastructure 12c Release 2 (12.2) Grid home as the active Oracle Clusterware home:

- a. On any of the cluster member nodes where the `rootupgrade.sh` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.
- b. Use the following command to start the installer, where `/u01/app/12.2.0/grid` is the location of the new (upgraded) Grid home:

```
cd /u01/app/12.2.0/grid/oui/bin
./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=false
ORACLE_HOME=/u01/app/12.2.0/grid
"CLUSTER_NODES=node1,node2,node3"
-doNotUpdateNodeList
```

Add the flag `-cfs` if the Grid home is a shared home.

5. Set Oracle Grid Infrastructure 12c Release 1 (12.1) Grid home as the active Oracle Clusterware home:

- a. On any of the cluster member nodes where the `rootupgrade` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.
- b. Use the following command to start the installer, where the path you provide for `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation.

```
$ cd /u01/app/12.1.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=true
ORACLE_HOME=/u01/app/12.1.0/grid
"CLUSTER_NODES=node1,node2,node3"
```

6. Start the 12.1 Oracle Clusterware stack on all nodes.

```
crsctl start crs
```

7. On any node, remove the MGMTDB resource as follows:

```
121_Grid_home/bin/srvctl remove mgmtdb
```

8. If you are downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1.0.2), run the following commands to configure the Grid Infrastructure Management Database:

- a. Run DBCA in the silent mode from the 12.1.0.2 Oracle home and create the Management Database container database (CDB) as follows:

```
12102_Grid_home/bin/dbca -silent -createDatabase -createAsContainerDatabase
true
-templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtdb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation 12102_Grid_home/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords -skipUserTemplateCheck
```

- b. Run DBCA in the silent mode from the 12.1.0.2 Oracle home and create the Management Database pluggable database (PDB) as follows:

```
12102_Grid_home/bin/dbca -silent -createPluggableDatabase -sourceDB
-MGMTDB -pdbName cluster_name -createPDBFrom RMANBACKUP
-PDBBackUpfile 12102_Grid_home/assistants/dbca/templates/mgmtseed_pdb.dfb
```

```
-PDBMetadataFile 12102_Grid_home/assistants/dbca/templates/mgmtseed_pdb.xml
-createAsClone true -internalSkipGIHomeCheck
```

9. If you are downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1.0.1), run DBCA in the silent mode from the 12.1.0.1 Oracle home and create the Management Database as follows:

```
12101_Grid_home/bin/dbca -silent -createDatabase
-templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtpdb
-storageType ASM -diskGroupName ASM_DG_NAME
-datafileJarLocation 12101_Grid_home/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords
```

10. Configure the Management Database by running the Configuration Assistant from the location `121_Grid_home/bin/mgmtca`.

11.13.4 Downgrading to Oracle Grid Infrastructure 11g Release 2 (11.2)

Use this procedure to downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).

1. Delete the Oracle Grid Infrastructure 12c Release 2 (12.2) Management Database:

```
dbca -silent -deleteDatabase -sourceDB -MGMTDB
```

2. Use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade` to stop the Oracle Grid Infrastructure 12c Release 2 (12.2) resources, and to shut down the stack. Run this command from a directory that has write permissions for the Oracle Grid Infrastructure installation user.

You can run the downgrade script in parallel on all cluster nodes, but one.

3. Downgrade the last node after you downgrade all other nodes:

```
# /u01/app/12.2.0/grid/crs/install/rootcrs.sh -downgrade
```

4. Follow these steps to remove Oracle Grid Infrastructure 12c Release 2 (12.2) Grid home as the active Oracle Clusterware home:

- a. On any of the cluster member nodes where the `rootupgrade.sh` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.
- b. Use the following command to start the installer, where `/u01/app/12.2.0/grid` is the location of the new (upgraded) Grid home:

```
$ cd /u01/app/12.2.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -updateNodeList
-silent CRS=false ORACLE_HOME=/u01/app/12.2.0/grid
"CLUSTER_NODES=node1,node2,node3" -doNotUpdateNodeList
```

Add the `-cfs` option if the Grid home is a shared home.

5. Follow these steps to set the Oracle Grid Infrastructure 11g Release 2 (11.2) Grid home as the active Oracle Clusterware home:

- a. On any of the cluster member nodes where the `rootupgrade` script has run successfully, log in as the Oracle Grid Infrastructure installation owner.
- b. Use the following command to start the installer, where the path you provide for the `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation.

```
$ cd /u01/app/11.2.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -
updateNodeList
-silent CRS=true ORACLE_HOME=/u01/app/11.2.0/grid
```

Add the `-cfs` option if the Grid home is a shared home.

6. Start the Oracle Clusterware stack manually from the earlier release Oracle Clusterware home using the command `crsctl start crs`. For example, where the earlier release home is `/u01/app/11.2.0/grid`, use the following command on each node:

```
/u01/app/11.2.0/grid/bin/crsctl start crs
```

11.13.5 Downgrading Oracle Grid Infrastructure after Upgrade Fails

If upgrade of Oracle Grid Infrastructure fails before setting the active version of Oracle Clusterware, then follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

Run this procedure to downgrade Oracle Clusterware only when the upgrade fails before root script runs the `crsctl set crs activeversion` command on the last node. Use this procedure for downgrading Oracle Grid Infrastructure if there is a need to avoid downtime of the whole cluster. This procedure downgrades the cluster to the previous release. Because Oracle ASM and database operations are limited in this state, it is recommended to move the cluster from this state as soon as possible. Complete the downgrade of Oracle Grid Infrastructure as per the procedure documented in *Downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1)* or *Downgrading to Oracle Grid Infrastructure 11g Release 2 (11.2)* for your software release.

1. Shut down the Oracle Grid Infrastructure stack on the first node:

```
crsctl stop crs
```

2. From any node where the Grid Infrastructure stack from the earlier release is running, unset the Oracle ASM rolling migration mode as follows:
 - a. Log in as `grid` user, and run the following command as `SYSASM` user on the Oracle ASM instance:

```
SQL> ALTER SYSTEM STOP ROLLING MIGRATION;
```

3. If you are upgrading from 11.2.0.4 or 12.1.0.1, then apply the latest available patches on all nodes in the cluster. If the pre-upgrade version is 12.1.0.2 or later, then patch is not required.
 - a. On all other nodes except the first node, where the earlier release Grid Infrastructure stack is running, apply the latest patch using the `opatch auto` procedure.
 - b. On the first node where the earlier release Grid Infrastructure stack is stopped, apply the latest patch using the `opatch apply` procedure.

For the list of latest available patches, see My Oracle Support at the following link:

<https://support.oracle.com/>

- i. Unlock the Grid Infrastructure home from the earlier release:

```
rootcrs.pl -unlock -crshome pre-upgrade-grid-home
```

pre-upgrade-grid-home is the previous release Grid home.

- ii. Apply the patch:

```
opatch apply -local -oh pre-upgrade-grid-home
```

- iii. Relock the Grid home from the earlier release:

```
rootcrs.pl -lock
```

- c. From any other node where the Grid Infrastructure stack from the earlier release is running, unset the Oracle ASM rolling migration mode as explained in step 2.

4. On any node running Oracle Grid Infrastructure other than the first node, from the Grid home of the earlier release, run the command:

```
clscfg -nodedowngrade -h hostname
```

hostname is the host name of the first node.

5. From the later release Grid home, run the command to downgrade Oracle Clusterware:

```
rootcrs.sh -downgrade -online
```

If `rootcrs.sh` is not present, then use `rootcrs.pl`.

6. Start Oracle Grid Infrastructure stack on the first node from the earlier release Grid home:

```
crsctl start crs
```

Note: You can downgrade the cluster nodes in any sequence.

Related Topics:

[Downgrading to Oracle Grid Infrastructure 12c Release 1 \(12.1\)](#) (page 11-24)

[Downgrading to Oracle Grid Infrastructure 11g Release 2 \(11.2\)](#) (page 11-26)

11.13.6 Downgrading Oracle Grid Infrastructure after Upgrade Fails on Remote Nodes

If upgrade of Oracle Grid Infrastructure fails on remote nodes, then you can follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

Run this procedure from an already upgraded node where the Grid Infrastructure stack of the latest release is running, and downgrade each node where upgrade has failed or completed.

1. Shut down the Oracle Grid Infrastructure stack on the remote node that is being downgraded:

```
crsctl stop crs
```

2. From the later release Grid home, run the command:

```
clscfg -nodedowngrade -h hostname
```

hostname is the host name of the remote node that is being downgraded.

3. From the remote node that is being downgraded, run the command:

```
rootcrs.pl -downgrade -online
```

4. Start Oracle Grid Infrastructure stack on the remote node from the earlier release Grid home:

```
crsctl start crs
```

5. After all remote nodes are downgraded, downgrade the last remaining node using the procedure described in *Downgrading Oracle Grid Infrastructure after Upgrade Fails*.

Note: You can downgrade the cluster nodes in any sequence.

Related Topics:

[Downgrading Oracle Grid Infrastructure after Upgrade Fails](#) (page 11-27)

If upgrade of Oracle Grid Infrastructure fails before setting the active version of Oracle Clusterware, then follow these steps to downgrade Oracle Grid Infrastructure to the earlier release.

11.14 Completing Failed or Interrupted Installations and Upgrades

If Oracle Universal Installer (OUI) exits on the node from which you started the upgrade, or the node reboots before you confirm that the `rootupgrade.sh` script was run on all nodes, then the upgrade remains incomplete.

In an incomplete installation or upgrade, configuration assistants still need to run, and the new Grid home still needs to be marked as active in the central Oracle inventory. You must complete the installation or upgrade on the affected nodes manually.

[Completing Failed Installations and Upgrades](#) (page 11-29)

Understand how to join nodes to the cluster after installation or upgrade fails on some nodes.

[Continuing Incomplete Upgrade of First Nodes](#) (page 11-30)

Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

[Continuing Incomplete Upgrades on Remote Nodes](#) (page 11-30)

Review this information to continue incomplete upgrade on remote nodes.

[Continuing Incomplete Installation on First Node](#) (page 11-31)

Review this information to continue an incomplete installation of Oracle Grid Infrastructure, if installation fails on the first node.

[Continuing Incomplete Installation on Remote Nodes](#) (page 11-31)

Review this information to continue incomplete installation on remote nodes.

11.14.1 Completing Failed Installations and Upgrades

Understand how to join nodes to the cluster after installation or upgrade fails on some nodes.

If installation or upgrade of Oracle Grid Infrastructure on some nodes fails, then the installation or upgrade completes with only successful nodes in the cluster. Follow this procedure to add the failed nodes to the cluster.

1. Remove the Oracle Grid Infrastructure software from the failed nodes:

```
Grid_home/deinstall/deinstall -local
```

2. As root user, from a node where Oracle Clusterware is installed, delete the failed nodes using the `delete node` command:

```
Grid_home/bin/crsctl delete node -n node_name
```

node_name is the node to be deleted.

3. Run the Oracle Grid Infrastructure installation wizard and follow the steps in the wizard to add the nodes:

```
Grid_home/gridSetup.sh
```

Alternatively, you can also add the nodes by running the `addnode` script:

```
Grid_home/addnode/addnode.sh
```

The nodes are added to the cluster.

11.14.2 Continuing Incomplete Upgrade of First Nodes

Review this information to complete the upgrade, if upgrade of Oracle Grid Infrastructure fails on the first node.

1. If the root script failure indicated a need to reboot, through the message `CLSRSC-400`, then reboot the first node (the node where the upgrade was started). Otherwise, manually fix or clear the error condition, as reported in the error output.
2. If necessary, log in as `root` to the first node. Change directory to the new Grid home on the first node, and run the `rootupgrade.sh` script on that node again. For example:

```
[root@node1]# cd /u01/app/12.2.0/grid  
[root@node1]# ./rootupgrade.sh
```

3. Complete the upgrade of all other nodes in the cluster.

```
[root@node1]# ./rootupgrade.sh
```

4. Configure a response file, and provide passwords for the installation.
5. To complete the upgrade, log in as the Grid installation owner, and run `gridSetup.sh`, located in the `Grid_home`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ gridSetup.sh -executeConfigTools -responseFile  
Grid_home/install/response/gridinstall.rsp
```

11.14.3 Continuing Incomplete Upgrades on Remote Nodes

Review this information to continue incomplete upgrade on remote nodes.

1. If the root script failure indicated a need to reboot, through the message CLSRSC-400, then reboot the first node (the node where the upgrade was started). Otherwise, manually fix or clear the error condition, as reported in the error output.
2. If root automation is being used, click **Retry** on the OUI instance on the first node.
3. If root automation is not being used, log into the affected node as `root`. Change directory to the Grid home, and run the `rootupgrade.sh` script on that node. For example:

```
[root@node6]# cd /u01/app/12.2.0/grid
[root@node6]# ./rootupgrade.sh
```

11.14.4 Continuing Incomplete Installation on First Node

Review this information to continue an incomplete installation of Oracle Grid Infrastructure, if installation fails on the first node.

1. If the root script failure indicated a need to reboot, through the message CLSRSC-400, then reboot the first node (the node where the installation was started). Otherwise, manually fix or clear the error condition, as reported in the error output.

2. If necessary, log in as `root` to the first node. Run the `oraInstRoot.sh` script on that node again. For example:

```
$ sudo -s
[root@node1]# cd /u01/app/oraInventory
[root@node1]# ./oraInstRoot.sh
```

3. Change directory to the Grid home on the first node, and run the `root` script on that node again. For example:

```
[root@node1]# cd /u01/app/12.1.0/grid
[root@node1]# ./root.sh
```

4. Complete the installation on all other nodes.
5. Configure a response file, and provide passwords for the installation.
6. To complete the installation, log in as the Grid installation owner, and run `gridSetup.sh`, located in the `Grid_home`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ gridSetup.sh -executeConfigTools -responseFile Oracle_home/install/
response/gridinstall.rsp
```

11.14.5 Continuing Incomplete Installation on Remote Nodes

Review this information to continue incomplete installation on remote nodes.

1. If the root script failure indicated a need to reboot, through the message CLSRSC-400, then reboot the affected node. Otherwise, manually fix or clear the error condition, as reported in the error output.
2. If root automation is being used, click **Retry** on the OUI instance on the first node.
3. If root automation is not being used, follow these steps:

- a. Log into the affected node as root, and run the `oraInstRoot.sh` script on that node. For example:

```
$ sudo -s
[root@node6]# cd /u01/app/oraInventory
[root@node6]# ./oraInstRoot.sh
```

- b. Change directory to the Grid home, and run the `root.sh` script on the affected node. For example:

```
[root@node6]# cd /u01/app/12.2.0/grid
[root@node6]# ./root.sh
```

4. Continue the installation from the OUI instance on the first node.

11.15 Converting to Oracle Extended Cluster After Upgrading Oracle Grid Infrastructure

Review this information to convert to an Oracle Extended Cluster after upgrading Oracle Grid Infrastructure. Oracle Extended Cluster enables you to deploy Oracle RAC databases on a cluster, in which some of the nodes are located in different sites.

Ensure that you have upgraded to Oracle Grid Infrastructure 12c Release 2 (12.2) as described in this chapter.

1. As root user, log in to the first node, and run the command:

```
rootcrs.sh -converttoextended -first -sites list_of_sites -site node_site
```

list_of_sites is the comma-separated list of sites in the extended cluster, and *node_site* is the node containing the site.

For example:

```
rootcrs.sh -converttoextended -first -sites newyork,newjersey,conn -site newyork
```

2. As root user, on all other nodes, run the following command:

```
rootcrs.sh -converttoextended -site node_site
```

node_site is the node containing the site.

For example:

```
rootcrs.sh -converttoextended -site newjersey
```

3. Optional: Delete the default site after the associated nodes and storage are migrated.

```
crsctl delete cluster site site_name
```

For example:

```
[root@node4]#crsctl delete cluster site mycluster
```

Removing Oracle Database Software

These topics describe how to remove Oracle software and configuration files.

You can remove Oracle software in one of two ways: Use Oracle Universal Installer with the `deinstall` option, or use the deinstallation tool (`deinstall`) that is included in Oracle homes. Oracle does not support the removal of individual products or components.

Caution:

If you have a standalone database on a node in a cluster, and if you have multiple databases with the same global database name (GDN), then you cannot use the `deinstall` tool to remove one database only.

[About Oracle Deinstallation Options](#) (page 12-2)

You can stop and remove Oracle Database software and components in an Oracle Database home with Oracle Universal Installer.

[Oracle Deinstallation Tool \(Deinstall\)](#) (page 12-4)

The `deinstall` tool is a script that you can run separately from Oracle Universal Installer (OUI).

[Deinstallation Examples for Oracle Database](#) (page 12-6)

Use these examples to help you understand how to run deinstallation using OUI (`runinstaller`) or as a standalone tool (`deinstall`).

[Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster](#) (page 12-7)

You can run the deinstallation tool with the `-paramfile` option to use the values you specify in the response file.

[Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster](#) (page 12-10)

If you have an Oracle Database installation using Oracle Restart (that is, an Oracle Grid Infrastructure installation for a standalone server), and you want to configure that server as a cluster member node, then complete the following tasks:

[Relinking Oracle Grid Infrastructure for a Cluster Binaries](#) (page 12-12)

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to modify the binaries, then use the following procedure, where `Grid_home` is the Oracle Grid Infrastructure for a cluster home:

[Changing the Oracle Grid Infrastructure Home Path](#) (page 12-12)

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to

change the Grid home path, then use the following example as a guide to detach the existing Grid home, and to attach a new Grid home:

[Unconfiguring Oracle Clusterware Without Removing Binaries](#) (page 12-13)
Running the `rootcrs.sh` command flags `-deconfig -force` enables you to unconfigure Oracle Clusterware on one or more nodes without removing installed binaries.

[Unconfiguring Oracle Member Cluster](#) (page 12-14)
Run this procedure to unconfigure Oracle Member Cluster.

12.1 About Oracle Deinstallation Options

You can stop and remove Oracle Database software and components in an Oracle Database home with Oracle Universal Installer.

You can remove the following software using Oracle Universal Installer or the Oracle deinstallation tool:

- Oracle Database
- Oracle Grid Infrastructure, which includes Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM)
- Oracle Real Application Clusters (Oracle RAC)
- Oracle Database Client

Starting with Oracle Database 12c, the deinstallation tool is integrated with the database installation media. You can run the deinstallation tool using the `runInstaller` command with the `-deinstall` and `-home` options from the base directory of the Oracle Database or Oracle Database Client installation media.

The deinstallation tool is also available as a separate command (`deinstall`) in Oracle home directories after installation. It is located in the `$ORACLE_HOME/deinstall` directory.

The deinstallation tool creates a response file by using information in the Oracle home and using the information you provide. You can use a response file that you generated previously by running the `deinstall` command using the `-checkonly` option. You can also edit the response file template.

If you run the deinstallation tool to remove an Oracle Grid Infrastructure installation, then the deinstaller prompts you to run the `deinstall` script as the `root` user. For Oracle Grid Infrastructure for a cluster, the script is `rootcrs.sh`, and for Oracle Grid Infrastructure for a standalone server (Oracle Restart), the script is `roothas.sh`.

Note:

- You must run the deinstallation tool from the same release to remove Oracle software. Do not run the deinstallation tool from a later release to remove Oracle software from an earlier release. For example, do not run the deinstallation tool from the 12.2 installation media to remove Oracle software from an existing 11.2.0.4 Oracle home.
 - Starting with Oracle Database 12c Release 1 (12.1.0.2), the `roothas.sh` script replaces the `roothas.pl` script in the Oracle Grid Infrastructure home for Oracle Restart, and the `rootcrs.sh` script replaces the `rootcrs.pl` script in the Grid home for Oracle Grid Infrastructure for a cluster.
-
-

If the software in the Oracle home is not running (for example, after an unsuccessful installation), then the deinstallation tool cannot determine the configuration, and you must provide all the configuration details either interactively or in a response file.

In addition, before you run the deinstallation tool for Oracle Grid Infrastructure installations:

- Dismount Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and disable Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM).
- If Grid Naming Service (GNS) is in use, then notify your DNS administrator to delete the subdomain entry from the DNS.

Files Deleted by the Deinstallation Tool

When you run the deinstallation tool, if the central inventory (`oraInventory`) contains no other registered homes besides the home that you are deconfiguring and removing, then the `deinstall` command removes the following files and directory contents in the Oracle base directory of the Oracle Database installation owner:

- `admin`
- `cfgtoollogs`
- `checkpoints`
- `diag`
- `oradata`
- `fast_recovery_area`

Oracle strongly recommends that you configure your installations using an Optimal Flexible Architecture (OFA) configuration, and that you reserve Oracle base and Oracle home paths for exclusive use of Oracle software. If you have any user data in these locations in the Oracle base that is owned by the user account that owns the Oracle software, then the deinstallation tool deletes this data.

Caution: The deinstallation tool deletes Oracle Database configuration files, user data, and fast recovery area (FRA) files even if they are located outside of the Oracle base directory path.

12.2 Oracle Deinstallation Tool (Deinstall)

The `deinstall` tool is a script that you can run separately from Oracle Universal Installer (OUI).

Purpose

The `deinstall` tool stops Oracle software, and removes Oracle software and configuration files on the operating system for a specific Oracle home.

Syntax

The standalone deinstallation tool uses the following syntax:

```
./deinstall [-silent] [-checkonly] [-paramfile complete path of input response file]
[-params name1=value name2=value . . .]
[-o complete path of directory for saving files]
[-tmpdir complete path of temporary directory to use]
[-logdir complete path of log directory to use] [-skipLocalHomeDeletion] [-skipRemoteHomeDeletion] [-help]
```

The `deinstall` tool run as a command option from OUI uses the following syntax, where *path* is the complete path to the home or file you specify:

```
./runInstaller -deinstall -home path [-silent] [-checkonly]
[-paramfile path] [-params name1=value name2=value . . .]
[-o path] [-tmpdir complete path of temporary directory to use]
[-logdir complete path of log directory to use] [-skipLocalHomeDeletion] [-skipRemoteHomeDeletion] [-help]
```

Parameters

Parameter	Description
-home	<p>Use this flag to indicate the home path of the Oracle home to check or deinstall.</p> <p>To deinstall Oracle software using the <code>deinstall</code> command, located in the Oracle home you plan to deinstall, provide a response file located outside the Oracle home, and do not use the <code>-home</code> flag.</p> <p>If you run the deinstallation tool from the <code>\$ORACLE_HOME/deinstall</code> path, then the <code>-home</code> flag is not required because the tool identifies the location of the home where it is run. If you use <code>runInstaller -deinstall</code> from the installation media, then <code>-home</code> is mandatory.</p>

Parameter	Description
<code>-silent</code>	<p>Use this flag to run the deinstallation tool in noninteractive mode. This option requires one of the following:</p> <ul style="list-style-type: none"> • A working system that it can access to determine the installation and configuration information. The <code>-silent</code> flag does not work with failed installations. • A response file that contains the configuration values for the Oracle home that is being deinstalled or deconfigured. <p>You can generate a response file to use or modify by running the tool with the <code>-checkonly</code> flag. The tool then discovers information from the Oracle home to deinstall and deconfigure. It generates the response file that you can then use with the <code>-silent</code> option.</p> <p>You can also modify the template file <code>deinstall.rsp.tmpl</code>, located in the <code>\$ORACLE_HOME/deinstall/response</code> directory.</p>
<code>-checkonly</code>	<p>Use this flag to check the status of the Oracle software home configuration. Running the deinstallation tool with the <code>-checkonly</code> flag does not remove the Oracle configuration. The <code>-checkonly</code> flag generates a response file that you can use with the deinstallation tool and <code>-silent</code> option.</p>
<code>-paramfile</code> <i>complete path of input response file</i>	<p>Use this flag to run the deinstallation tool with a response file in a location other than the default. When you use this flag, provide the complete path where the response file is located.</p> <p>The default location of the response file depends on the location of the deinstallation tool:</p> <ul style="list-style-type: none"> • From the installation media or stage location: <code>/response</code> • After installation from the installed Oracle home: <code>\$ORACLE_HOME/deinstall/response</code>
<code>-params</code> [<i>name1=value name2=value name3=value . . .</i>]	<p>Use this flag with a response file to override one or more values to change in a response file you have created.</p>

Parameter	Description
<code>-o</code> <i>complete path of directory for saving response files</i>	Use this flag to provide a path other than the default location where the response file (<code>deinstall.rsp.tmp1</code>) is saved. The default location of the response file depends on the location of the deinstallation tool: <ul style="list-style-type: none"> • From the installation media or stage location: <code>/response</code> • After installation from the installed Oracle home: <code>\$ORACLE_HOME/deinstall/response</code>
<code>-tmpdir</code> <i>complete path of temporary directory to use</i>	Use this flag to specify a non-default location where Oracle Deinstallation Tool writes the temporary files for the deinstallation.
<code>-logdir</code> <i>complete path of log directory to use</i>	Use this flag to specify a non-default location where Oracle Deinstallation Tool writes the log files for the deinstallation.
<code>-local</code>	Use this flag on a multinode environment to deinstall Oracle software in a cluster. When you run <code>deinstall</code> with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where <code>deinstall</code> is run). On remote nodes, it deconfigures Oracle software, but does not deinstall the Oracle software.
<code>-skipLocalHomeDeletion</code>	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a local Grid home without deleting the Grid home.
<code>-skipRemoteHomeDeletion</code>	Use this flag in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a remote Grid home without deleting the Grid home.
<code>-help</code>	Use this option to obtain additional information about the command option flags.

12.3 Deinstallation Examples for Oracle Database

Use these examples to help you understand how to run deinstallation using OUI (`runinstaller`) or as a standalone tool (`deinstall`).

If you run the deinstallation tool from the installation media using `runInstaller -deinstall`, then help is displayed that guides you through the deinstallation process. You can also use the `-home` flag and provide a path to the home directory of the Oracle software to remove from your system. If you have a response file, then use the optional flag `-paramfile` to provide a path to the response file.

You can generate a deinstallation response file by running the deinstallation tool with the `-checkonly` flag. Alternatively, you can use the response file template located at `$ORACLE_HOME/deinstall/response/deinstall.rsp.tmpl`.

In the following example, the `runInstaller` command is in the path `/directory_path`, where `/directory_path` is the path to the database directory on the installation media, and `/u01/app/oracle/product/12.2.0/dbhome_1/` is the path to the Oracle home you want to remove:

```
$ cd /directory_path/
$ ./runInstaller -deinstall -home /u01/app/oracle/product/12.2.0/dbhome_1/
```

The following example uses a response file called `my_db_paramfile.tmpl` in the software owner location `/home/usr/oracle`:

```
$ cd /directory_path/
$ ./runInstaller -deinstall -paramfile /home/usr/oracle/my_db_paramfile.tmpl
```

If you run the deinstallation tool using `deinstall` from the `$ORACLE_HOME/deinstall` directory, then the deinstallation starts without prompting you for the Oracle home path.

In the following example, the `deinstall` command is in the path `/u01/app/oracle/product/12.2.0/dbhome_1/deinstall`. It uses a response file called `my_db_paramfile.tmpl` in the software owner location `/home/usr/oracle`:

```
$ cd /u01/app/oracle/product/12.2.0/dbhome_1/deinstall
$ ./deinstall -paramfile /home/usr/oracle/my_db_paramfile.tmpl
```

To remove the Oracle Grid Infrastructure home, use the deinstallation script in the Oracle Grid Infrastructure home.

In this example, the Oracle Grid Infrastructure home is `/u01/app/oracle/product/12.2.0/grid`:

```
$ cd /u01/app/oracle/product/12.2.0/grid/deinstall$ ./deinstall -paramfile /home/usr/oracle/my_grid_paramfile.tmpl
```

12.4 Deinstallation Response File Example for Oracle Grid Infrastructure for a Cluster

You can run the deinstallation tool with the `-paramfile` option to use the values you specify in the response file.

The following is an example of a response file for a cluster on nodes `node1` and `node2`, in which the Oracle Grid Infrastructure for a cluster software binary owner is `grid`, the Oracle Grid Infrastructure home (Grid home) is in the path `/u01/app/12.2.0/grid`, the Oracle base (the Oracle base for Oracle Grid Infrastructure, containing Oracle ASM log files, Oracle Clusterware logs, and other administrative files) is `/u01/app/grid/`, the central Oracle Inventory home (`oraInventory`) is `/u01/app/oraInventory`, the virtual IP addresses (VIP) are `192.0.2.2` and `192.0.2.4`, the local node (the node where you run the deinstallation session from) is `node1`:

```
# Copyright (c) 2005, 2016 Oracle Corporation. All rights reserved.
ORACLE_HOME=/u01/app/12.2.0/grid
CDATA_AUSIZE=4
BIG_CLUSTER=true
ISROLLING=true
LOCAL_NODE=node1
OCR_VD_DISKGROUPS="+DATA1"
```

```
MGMTDB_DIAG=/u01/app/grid
OCRID=
MGMTDB_SPFIL=" +DATA1/_MGMTDB/PARAMETERFILE/spfile.271.923210081"
ObaseCleanupPtrLoc=/tmp/deinstall2016-10-06_09-36-04AM/utl/orabase_cleanup.lst
CDATA_BACKUP_QUORUM_GROUPS=
ASM_CREDENTIALS=
MGMTDB_NODE_LIST=node1,node2
EXTENDED_CLUSTER=false
LISTENER_USERNAME=cuser
local=false
inventory_loc=/u01/app/oraInventory
ORACLE_HOME=/u01/app/12.2.0/grid
ASM_HOME=/u01/app/grid
ASM_DISK_GROUPS="+DATA1"
HUB_NODE_VIPS=AUTO,AUTO
PING_TARGETS=
ORA_DBA_GROUP=oinstall
ASM_DISCOVERY_STRING=/dev/rdsk/*
CDATA_DISKS=/dev/rdsk/c0t600144F0C4A01A3F000056E6A12A0022d0s3
MinimumSupportedVersion=11.2.0.1.0
NEW_HOST_NAME_LIST=
ORACLE_HOME_VERSION=12.2.0.1.0
PRIVATE_NAME_LIST=
MGMTDB_DB_UNIQUE_NAME=_mgmtdb
ASM_DISKSTRING=/dev/rdsk/* ,AFD:*
CDATA_QUORUM_GROUPS=
CRS_HOME=true
ODA_CONFIG=
JLIBDIR=/u01/app/jlib
CRFHOME="/u01/app/"
USER_IGNORED_PREREQ=true
MGMTDB_ORACLE_BASE=/u01/app/grid/
DROP_MGMTDB=true
RHP_CONF=false
OCRLOC=
GNS_TYPE=local
CRS_STORAGE_OPTION=1
CDATA_SITES=
GIMR_CONFIG=local
CDATA_BACKUP_SIZE=0
GPNPGCONFIGDIR=$ORACLE_HOME
MGMTDB_IN_HOME=true
CDATA_DISK_GROUP="+DATA2"
LANGUAGE_ID=AMERICAN_AMERICA.AL32UTF8
CDATA_BACKUP_FAILURE_GROUPS=
CRS_NODEVIPS='AUTO/255.255.254.0/net0,AUTO/255.255.254.0/net0'
ORACLE_OWNER=cuser
GNS_ALLOW_NET_LIST=
silent=true
INSTALL_NODE=node1.example.com
ORACLE_HOME_VERSION_VALID=true
inst_group=oinstall
LOGDIR=/tmp/deinstall2016-10-06_09-36-04AM/logs/
EXTENDED_CLUSTER_SITES=
CDATA_REDUNDANCY=EXTERNAL
CDATA_BACKUP_DISK_GROUP="+DATA2"
APPLICATION_VIP=
HUB_NODE_LIST=node1,node2
NODE_NAME_LIST=node1,node2
GNS_DENY_ITF_LIST=
ORA_CRS_HOME=/u01/app/12.2.0/grid/
```

```
JREDIR=/u01/app/12.2.0/grid/jdk/jre/
ASM_LOCAL_SID=+ASM1
ORACLE_BASE=/u01/app/
GNS_CONF=true
CLUSTER_CLASS=DOMAINSERVICES
ORACLE_BINARY_OK=true
CDATA_BACKUP_REDUNDANCY=EXTERNAL
CDATA_FAILURE_GROUPS=
ASM_CONFIG=near
OCR_LOCATIONS=
ASM_ORACLE_BASE=/u01/app/12.2.0/
OLRLOC=
GIMR_CREDENTIALS=
GPNPCONFIGDIR=$ORACLE_HOME
ORA_ASM_GROUP=asmadmin
GNS_CREDENTIALS=
CDATA_BACKUP_AUSIZE=4
GNS_DENY_NET_LIST=
OLD_CRS_HOME=
NEW_NODE_NAME_LIST=
GNS_DOMAIN_LIST=node1.example.com
ASM_UPGRADE=false
NETCA_LISTENERS_REGISTERED_WITH_CRS=LISTENER
CDATA_BACKUP_DISKS=/dev/rdsk/
ASMCA_ARGS=
CLUSTER_GUID=
CLUSTER_NODES=node1,node2
MGMTDB_NODE=node2
ASM_DIAGNOSTIC_DEST=/u01/app/
NEW_PRIVATE_NAME_LIST=
AFD_LABELS_NO_DG=
AFD_CONFIGURED=true
CLSCFG_MISSCOUNT=
MGMT_DB=true
SCAN_PORT=1521
ASM_DROP_DISKGROUPS=true
OPC_NAT_ADDRESS=
CLUSTER_TYPE=DB
NETWORKS="net0"/IP_Address:public,"net1"/IP_Address:asm,"net1"/
IP_Address:cluster_interconnect
OCR_VOTINGDISK_IN_ASM=true
HUB_SIZE=32
CDATA_BACKUP_SITES=
CDATA_SIZE=0
REUSEDG=false
MGMTDB_DATAFILE=
ASM_IN_HOME=true
HOME_TYPE=CRS
MGMTDB_SID="-MGMTDB"
GNS_ADDR_LIST=mycluster-gns.example.com
CLUSTER_NAME=node1-cluster
AFD_CONF=true
MGMTDB_PWDFILE=
OPC_CLUSTER_TYPE=
VOTING_DISKS=
SILENT=false
VNDR_CLUSTER=false
TZ=localtime
GPNP_PA=
DC_HOME=/tmp/deinstall2016-10-06_09-36-04AM/logs/
CSS_LEASEDURATION=400
```

```

REMOTE_NODES=node2
ASM_SPFILE=
NEW_NODEVIPS='n1-vip/255.255.252.0/eth0,n2-vip/255.255.252.0/eth0'
SCAN_NAME=node1-cluster-scan.node1-cluster.com
RIM_NODE_LIST=
INVENTORY_LOCATION=/u01/app/oraInventory

```

Note:

Do not use quotation marks with variables except in the following cases:

- Around addresses in CRS_NODEVIPS:

```
CRS_NODEVIPS='n1-vip/255.255.252.0/eth0,n2-vip/255.255.252.0/eth0'
```

- Around interface names in NETWORKS:

```
NETWORKS="eth0"/192.0.2.1\:\:public,"eth1"/10.0.0.1\:\:cluster_interconnect
VIP1_IP=192.0.2.2
```

12.5 Migrating Standalone Oracle Grid Infrastructure Servers to a Cluster

If you have an Oracle Database installation using Oracle Restart (that is, an Oracle Grid Infrastructure installation for a standalone server), and you want to configure that server as a cluster member node, then complete the following tasks:

1. Inspect the Oracle Restart configuration with `srvctl` using the following syntax, where *db_unique_name* is the unique name for the database, and *lsnrname* is the name of the listener:

```

srvctl config database -db db_unique_name
srvctl config service -db db_unique_name
srvctl config listener -listener lsnrname

```

Write down the configuration information for the server.

2. Stop all of the databases, services, and listeners that you discovered in step 1.
3. If present, unmount all Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems.
4. Log in as `root`, and change directory to *Grid home/crs/install*. For example:

```
# cd /u01/app/12.2.0/grid/crs/install
```

5. Unconfigure the Oracle Grid Infrastructure installation for a standalone server (Oracle Restart), using the following command:

```
# roothas.sh -deconfig -force
```

6. Prepare the server for Oracle Clusterware configuration, as described in this document. In addition, you can install Oracle Grid Infrastructure for a cluster in the same location as Oracle Restart, or in a different location.

Installing in the Same Location as Oracle Restart

- a. Proceed to step 7.

Installing in a Different Location than Oracle Restart

- a. Set up Oracle Grid Infrastructure software in the new Grid home software location as described in *Installing Only the Oracle Grid Infrastructure Software*.
- b. Proceed to step 7.

7. Set the environment variables as follows:

```
export oracle_install_asm_UseExistingDG=true or false
export oracle_install_asm_DiskGroupName=disk_group_name
export oracle_install_asm_DiskDiscoveryString=asm_discovery_string
export oracle_install_asm_ConfigureGIMRDataDG=true or false
export oracle_install_asm_GIMRDataDGName=disk_group_name
```

8. As the Oracle Grid Infrastructure installation owner, run the installer.

You can complete the installation either interactively or in the silent mode. If you perform a silent installation, save and stage the response file as described in *Recording Response Files*.

9. After saving the response file, run the command:

```
$ Grid_home/gridSetup.sh -silent -responseFile $ORACLE_HOME/GI.rsp
```

10. Run `root.sh`.

11. Mount the Oracle ASM disk group used by Oracle Restart.

12. If you used Oracle ACFS with Oracle Restart, then:

- a. Start Oracle ASM Configuration Assistant (ASMCA). Run the `volenable` command to enable all Oracle Restart disk group volumes.
- b. Mount all Oracle ACFS file systems manually.

13. Add back Oracle Clusterware services to the Oracle Clusterware home, using the information you wrote down in step 1, including adding back Oracle ACFS resources. For example:

```
/u01/app/grid/product/12.2.0/grid/bin/srvctl add filesystem -device
/dev/asm/db1 -diskgroup ORestartData -volume db1 -mountpointpath
/u01/app/grid/product/12.2.0/db1 -user grid
```

14. Add the Oracle Database for support by Oracle Grid Infrastructure for a cluster, using the configuration information you recorded in step 1. Use the following command syntax, where *db_unique_name* is the unique name of the database on the node, and *nodename* is the name of the node:

```
srvctl add database -db db_unique_name -spfile -pwfile -
oraclehome $ORACLE_HOME -node nodename
```

- a. For example, first verify that the `ORACLE_HOME` environment variable is set to the location of the database home directory.

b. Next, to add the database name `mydb`, enter the following command:

```
srvctl add database -db mydb -spfile -pwfile -oraclehome $ORACLE_HOME -node
node1
```

- c. Add each service to the database, using the command `srvctl add service`. For example, add `myservice` as follows:

```
srvctl add service -db mydb -service myservice
```

15. Add nodes to your cluster, as required, using the Oracle Grid Infrastructure installer.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about adding nodes to your cluster.

12.6 Relinking Oracle Grid Infrastructure for a Cluster Binaries

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to modify the binaries, then use the following procedure, where *Grid_home* is the Oracle Grid Infrastructure for a cluster home:

Caution:

Before relinking executables, you must shut down all executables that run in the Oracle home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries. If present, unmount all Oracle Automatic Storage Management Cluster File System (Oracle ACFS) filesystems.

As root:

```
# cd Grid_home/crs/install
# rootcrs.sh -unlock
```

As the Oracle Grid Infrastructure for a cluster owner:

```
$ export ORACLE_HOME=Grid_home
$ Grid_home/bin/relink
```

As root again:

```
# cd Grid_home/rdbms/install/
# ./rootadd_rdbms.sh
# cd Grid_home/crs/install
# rootcrs.sh -lock
```

You must relink the Oracle Clusterware and Oracle ASM binaries every time you apply an operating system patch or after you perform an operating system upgrade that does not replace the root file system. For an operating system upgrade that results in a new root file system, you must remove the node from the cluster and add it back into the cluster.

For upgrades from previous releases, if you want to deinstall the prior release Grid home, then you must first unlock the prior release Grid home. Unlock the previous release Grid home by running the command `rootcrs.sh -unlock` from the previous release home. After the script has completed, you can run the deinstallation tool.

12.7 Changing the Oracle Grid Infrastructure Home Path

After installing Oracle Grid Infrastructure for a cluster (Oracle Clusterware and Oracle ASM configured for a cluster), if you need to change the Grid home path, then use the following example as a guide to detach the existing Grid home, and to attach a new Grid home:

Caution:

Before changing the Grid home, you must shut down all executables that run in the Grid home directory that you are relinking. In addition, shut down applications linked with Oracle shared libraries.

1. Log in as the Oracle Grid Infrastructure installation owner (grid).
2. Change directory to *Grid_home*/bin and, as root, run the command `crsctl stop crs`. For example:


```
$ cd /u01/app/12.2.0/grid/bin
$ ./crsctl stop crs
```
3. As grid user, detach the existing Grid home by running the following command, where `/u01/app/12.2.0/grid` is the existing Grid home location:


```
$ /u01/app/12.2.0/grid/oui/bin/runInstaller -silent -waitforcompletion\
-detachHome ORACLE_HOME='/u01/app/12.2.0/grid' -local
```
4. As root, move the Grid binaries from the old Grid home location to the new Grid home location. For example, where the old Grid home is `/u01/app/12.2.0/grid` and the new Grid home is `/u01/app/12c/`:


```
# mkdir /u01/app/12c
# cp -pR /u01/app/12.2.0/grid /u01/app/12c
```
5. Unlock the destination Grid home:


```
# cd /u01/app/12c/grid/crs/install
# rootcrs.sh -unlock -dstcrshome /u01/app/12c/grid
```
6. Clone the Oracle Grid Infrastructure installation, using the instructions provided in *Oracle Clusterware Administration and Deployment Guide*.

When you navigate to the *Grid_home/clone/bin* directory and run the `clone.pl` script, provide values for the input parameters that provide the path information for the new Grid home.

The Oracle Clusterware and Oracle ASM binaries are relinked when you clone the Oracle Grid Infrastructure installation.
7. As root again, enter the following command to start up in the new home location:


```
# cd /u01/app/12c/grid/crs/install
# rootcrs.sh -move -dstcrshome /u01/app/12c/grid
```
8. Repeat steps 1 through 8 on each cluster member node.

Caution: While cloning, ensure that you do not change the Oracle home base, otherwise the move operation fails.

12.8 Unconfiguring Oracle Clusterware Without Removing Binaries

Running the `rootcrs.sh` command flags `-deconfig -force` enables you to unconfigure Oracle Clusterware on one or more nodes without removing installed binaries.

This feature is useful if you encounter an error on one or more cluster nodes during installation when running the `root.sh` command, such as a missing operating system package on one node. By running `rootcrs.sh -deconfig -force` on nodes where you encounter an installation error, you can unconfigure Oracle Clusterware on those nodes, correct the cause of the error, and then run `root.sh` again.

Note:

Stop any databases, services, and listeners that may be installed and running before deconfiguring Oracle Clusterware. In addition, dismount Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and disable Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM) volumes.

Caution:

Commands used in this section remove the Oracle Grid infrastructure installation for the entire cluster. If you want to remove the installation from an individual node, then see *Oracle Clusterware Administration and Deployment Guide*.

To unconfigure Oracle Clusterware:

1. Log in as the `root` user on a node where you encountered an error.
2. Change directory to `Grid_home/crs/install`. For example:

```
# cd /u01/app/12.2.0/grid/crs/install
```
3. Run `rootcrs.sh` with the `-deconfig` and `-force` flags. For example:

```
# rootcrs.sh -deconfig -force
```

Repeat on other nodes as required.

4. If you are deconfiguring Oracle Clusterware on all nodes in the cluster, then on the last node, enter the following command:

```
# rootcrs.sh -deconfig -force -lastnode
```

The `-lastnode` flag completes deconfiguration of the cluster, including the OCR and voting files.

Caution:

Run the `rootcrs.sh -deconfig -force -lastnode` command on a Hub Node. Deconfigure all Leaf Nodes before you run the command with the `-lastnode` flag.

12.9 Unconfiguring Oracle Member Cluster

Run this procedure to unconfigure Oracle Member Cluster.

1. Run the deinstall tool to unconfigure Oracle Member Cluster:

```
Grid_home/deinstall/deinstall.sh
```

2. Complete the deinstallation by running the root script on all the nodes when prompted.

```
# rootcrs.sh -deconfig
```

3. Delete the Member Cluster Manifest File for the Oracle Member Cluster and stored on the Oracle Domain Services Cluster:

```
crsctl delete member_cluster_configuration member_cluster_name
```

Related Topics:

Oracle Clusterware Administration and Deployment Guide

Installing and Configuring Oracle Database Using Response Files

Review the following topics to install and configure Oracle products using response files.

[How Response Files Work](#) (page A-2)

Response files can assist you with installing an Oracle product multiple times on multiple computers.

[Reasons for Using Silent Mode or Response File Mode](#) (page A-2)

Review this section for use cases for running the installer in silent mode or response file mode.

[Using Response Files](#) (page A-3)

Review this information to use response files.

[Preparing Response Files](#) (page A-3)

Review this information to prepare response files for use during silent mode or response file mode installations.

[Running Oracle Universal Installer Using a Response File](#) (page A-6)

After creating the response file, run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation.

[Running Configuration Assistants Using Response Files](#) (page A-7)

You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on the system. To run configuration assistants in response file or silent mode, you must copy and edit a response file template.

[Postinstallation Configuration Using Response File Created During Installation](#) (page A-10)

Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.

[Postinstallation Configuration Using the ConfigToolAllCommands Script](#) (page A-12)

You can create and run a response file configuration after installing Oracle software. The `configToolAllCommands` script requires users to create a second response file, of a different format than the one used for installing the product.

A.1 How Response Files Work

Response files can assist you with installing an Oracle product multiple times on multiple computers.

When you start Oracle Universal Installer (OUI), you can use a response file to automate the installation and configuration of Oracle software, either fully or partially. OUI uses the values contained in the response file to provide answers to some or all installation prompts.

Typically, the installer runs in interactive mode, which means that it prompts you to provide information in graphical user interface (GUI) screens. When you use response files to provide this information, you run the installer from a command prompt using either of the following modes:

- Silent mode

If you include responses for all of the prompts in the response file and specify the `-silent` option when starting the installer, then it runs in silent mode. During a silent mode installation, the installer does not display any screens. Instead, it displays progress information in the terminal that you used to start it.

- Response file mode

If you include responses for some or all of the prompts in the response file and omit the `-silent` option, then the installer runs in response file mode. During a response file mode installation, the installer displays all the screens, screens for which you specify information in the response file, and also screens for which you did not specify the required information in the response file.

You define the settings for a silent or response file installation by entering values for the variables listed in the response file. For example, to specify the Oracle home name, provide the Oracle home path for the `ORACLE_HOME` environment variable:

```
ORACLE_HOME=/u01/app/oracle/product/12.2.0/dbhome_1
```

Another way of specifying the response file variable settings is to pass them as command-line arguments when you run the installer. For example:

```
-silent directory_path
```

In this command, *directory_path* is the path of the database directory on the installation media, or the path of the directory on the hard drive.

A.2 Reasons for Using Silent Mode or Response File Mode

Review this section for use cases for running the installer in silent mode or response file mode.

Mode	Uses
Silent	<p>Use silent mode for the following installations:</p> <ul style="list-style-type: none"> • Complete an unattended installation, which you schedule using operating system utilities such as <code>at</code>. • Complete several similar installations on multiple systems without user interaction. • Install the software on a system that does not have X Window System software installed on it. <p>The installer displays progress information on the terminal that you used to start it, but it does not display any of the installer screens.</p>
Response file	<p>Use response file mode to complete similar Oracle software installations on more than one system, providing default answers to some, but not all of the installer prompts.</p> <p>If you do not specify information required for a particular installer screen in the response file, then the installer displays that screen. It suppresses screens for which you have provided all of the required information.</p>

A.3 Using Response Files

Review this information to use response files.

Use the following general steps to install and configure Oracle products using the installer in silent or response file mode:

Note:

You must complete all required preinstallation tasks on a system before running the installer in silent or response file mode.

1. Create the `oraInst.loc` file if it is not present on the server.
2. Prepare a response file.
3. Run the installer in silent or response file mode.
4. Run the root scripts as prompted by Oracle Universal Installer.
5. If you completed a software-only installation, then run Net Configuration Assistant and DBCA in silent or response file mode.

A.4 Preparing Response Files

Review this information to prepare response files for use during silent mode or response file mode installations.

[Editing a Response File Template](#) (page A-4)

Oracle provides response file templates for each product and installation type, and for each configuration tool.

[Recording Response Files](#) (page A-5)

You can use OUI in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode

installations. This method is useful for Advanced or software-only installations.

A.4.1 Editing a Response File Template

Oracle provides response file templates for each product and installation type, and for each configuration tool.

For Oracle Database, the response file templates are located in the `database/response` directory on the installation media and in the `Oracle_home/inventory/response` directory after the software is installed. For Oracle Grid Infrastructure, the response file templates are located in the `Grid_home/install/response` directory after the software is installed.

Note:

If you copied the software to a hard disk, then the response files are located in the `/response` directory.

All response file templates contain comment entries, sample formats, examples, and other useful instructions. Read the response file instructions to understand how to specify values for the response file variables, so that you can customize your installation.

The following table lists the response files provided with this software:

Table A-1 Response Files for Oracle Database and Oracle Grid Infrastructure

Response File	Description
<code>db_install.rsp</code>	Silent installation of Oracle Database.
<code>dbca.rsp</code>	Silent creation and configuration of Oracle Database using Oracle DBCA.
<code>netca.rsp</code>	Silent configuration of Oracle Net using Oracle NETCA.
<code>grid_setup.rsp</code>	Silent configuration of Oracle Grid Infrastructure installations.

Caution:

When you modify a response file template and save a file for use, the response file may contain plain text passwords. Ownership of the response file should be given to the Oracle software installation owner only, and permissions on the response file should be changed to 600. Oracle strongly recommends that database administrators or other administrators delete or secure response files when they are not in use.

To copy and modify a response file:

1. Copy the response file from the response file directory to a directory on your system:

```
$ cp /Oracle_home/install/response/product_timestamp.rsp local_directory
```

2. Open the response file in a text editor:

```
$ vi /local_dir/response_file.rsp
```

3. Follow the instructions in the file to edit it.

Note:

The installer or configuration assistant fails if you do not correctly configure the response file. Also, ensure that your response file name has the `.rsp` suffix.

4. Secure the response file by changing the permissions on the file to 600:

```
$ chmod 600 /local_dir/response_file.rsp
```

Ensure that only the Oracle software owner user can view or modify response files or consider deleting them after the installation succeeds.

Note:

A fully-specified response file for an Oracle Database installation contains the passwords for database administrative accounts and for a user who is a member of the OSDBA group (required for automated backups).

Related Topics:

Oracle Universal Installer User's Guide

A.4.2 Recording Response Files

You can use OUI in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode installations. This method is useful for Advanced or software-only installations.

You can save all the installation steps into a response file during installation by clicking **Save Response File** on the Summary page. You can use the generated response file for a silent installation later.

When you record the response file, you can either complete the installation, or you can exit from the installer on the Summary page, before OUI starts to copy the software to the system.

If you use record mode during a response file mode installation, then the installer records the variable values that were specified in the original source response file into the new response file.

Note:

You cannot save passwords while recording the response file.

To record a response file:

1. Complete preinstallation tasks as for a standard installation.

When you run the installer to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required preinstallation tasks and record the response file while completing an installation.

2. Ensure that the Oracle software owner user (typically `oracle`) has permissions to create or write to the Oracle home path that you specify when you run the installer.
3. On each installation screen, specify the required information.
4. When the installer displays the Summary screen, perform the following steps:
 - a. Click **Save Response File**. In the window, specify a file name and location for the new response file. Click **Save** to write the responses you entered to the response file.
 - b. Click **Finish** to continue with the installation.Click **Cancel** if you do not want to continue with the installation. The installation stops, but the recorded response file is retained.

Note: Ensure that your response file name has the `.rsp` suffix.

5. If you do not complete the installation, then delete the Oracle home directory that the installer created using the path you specified in the Specify File Locations screen.
6. Before you use the saved response file on another system, edit the file and make any required changes. Use the instructions in the file as a guide when editing it.

A.5 Running Oracle Universal Installer Using a Response File

After creating the response file, run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation.

Run Oracle Universal Installer at the command line, specifying the response file you created. The Oracle Universal Installer executables, `runInstaller` and `gridSetup.sh`, provide several options. For help information on the full set of these options, run the `gridSetup.sh` or `runInstaller` command with the `-help` option. For example:

- For Oracle Database:

```
$ directory_path/runInstaller -help
```
- For Oracle Grid Infrastructure:

```
$ Grid_home/gridSetup.sh -help
```

The help information appears in a window after some time.

To run the installer using a response file:

1. Complete the preinstallation tasks for a normal installation.
2. Log in as the software installation owner user.
3. If you are completing a response file mode installation, then set the operating system `DISPLAY` environment variable for the user running the installation.

Note:

You do not have to set the `DISPLAY` environment variable if you are completing a silent mode installation.

4. To start the installer in silent or response file mode, enter a command similar to the following:

- For Oracle Database:

```
$ /directory_path/runInstaller [-silent] [-noconfig] \  
-responseFile responsefilename
```

- For Oracle Grid Infrastructure:

```
$ Grid_home/gridSetup.sh [-silent] [-noconfig] \  
-responseFile responsefilename
```

Note:

Do not specify a relative path to the response file. If you specify a relative path, then the installer fails.

In this example:

- `directory_path` is the path of the DVD or the path of the directory on the hard drive where you have copied the installation binaries.
 - `-silent` runs the installer in silent mode.
 - `-noconfig` suppresses running the configuration assistants during installation, and a software-only installation is performed instead.
 - `responsefilename` is the full path and file name of the installation response file that you configured.
5. When the installation completes, log in as the `root` user and run the `root.sh` script. For example

```
$ su root  
password:  
# /oracle_home_path/root.sh
```

6. If this is the first time you are installing Oracle software on your system, then Oracle Universal Installer prompts you to run the `oraInstRoot.sh` script.

Log in as the `root` user and run the `oraInstRoot.sh` script:

```
$ su root  
password:  
# /u01/app/oraInventory/oraInstRoot.sh
```

A.6 Running Configuration Assistants Using Response Files

You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on the system. To run configuration assistants in response file or silent mode, you must copy and edit a response file template.

Note:

If you copied the software to a hard disk, then the response file template is located in the `/response` directory.

Running Database Configuration Assistant Using Response Files (page A-8)

You can run Oracle Database Configuration Assistant (Oracle DBCA) in response file mode to configure and start an Oracle database on the system.

Running Net Configuration Assistant Using Response Files (page A-9)

You can run Net Configuration Assistant in silent mode to configure and start an Oracle Net Listener on the system, configure naming methods, and configure Oracle Net service names.

A.6.1 Running Database Configuration Assistant Using Response Files

You can run Oracle Database Configuration Assistant (Oracle DBCA) in response file mode to configure and start an Oracle database on the system.

To run Database Configuration Assistant in response file mode, you must copy and edit a response file template. Oracle provides a response file template named `dbca.rsp` in the `ORACLE_HOME/assistants/dbca` directory and also in the `/response` directory on the installation media. To run Oracle DBCA in response file mode, you must use the `-responseFile` flag in combination with the `-silent` flag. You must also use a graphical display and set the `DISPLAY` environment variable.

To run Database Configuration Assistant in response file mode:

1. Copy the `dbca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/dbca.rsp local_directory
```

In this example, `directory_path` is the path of the database directory on the DVD. If you have copied the software to a hard drive, you can edit the file in the response directory if you prefer.

As an alternative to editing the response file template, you can also create a database by specifying all required information as command line options when you run Oracle DBCA. For information about the list of options supported, enter the following command:

```
$ $ORACLE_HOME/bin/dbca -help
```

2. Open the response file in a text editor:

```
$ vi /local_dir/dbca.rsp
```

3. Follow the instructions in the file to edit the file.

Note:

Oracle DBCA fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.

5. To run Oracle DBCA in response file mode, set the `DISPLAY` environment variable.
6. Use the following command syntax to run Oracle DBCA in silent or response file mode using a response file:

```
$ORACLE_HOME/bin/dbca {-silent} -responseFile \  
/local_dir/dbca.rsp
```

In this example:

- `-silent` option indicates that Oracle DBCA runs in silent mode.
- `local_dir` is the full path of the directory where you copied the `dbca.rsp` response file template.

During configuration, Oracle DBCA displays a window that contains status messages and a progress bar.

A.6.2 Running Net Configuration Assistant Using Response Files

You can run Net Configuration Assistant in silent mode to configure and start an Oracle Net Listener on the system, configure naming methods, and configure Oracle Net service names.

To run Net Configuration Assistant in silent mode, you must copy and edit a response file template. Oracle provides a response file template named `netca.rsp` in the response directory in the `database/response` directory in the installation media.

Note:

If you copied the software to a hard disk, then the response file template is located in the `database/response` directory.

To run Net Configuration Assistant using a response file:

1. Copy the `netca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/netca.rsp local_directory
```

In this example, `directory_path` is the path of the `database` directory on the DVD. If you have copied the software to a hard drive, you can edit the file in the response directory if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/netca.rsp
```

3. Follow the instructions in the file to edit it.

Note:

Net Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.

5. Enter a command similar to the following to run Net Configuration Assistant in silent mode:

```
$ $ORACLE_HOME/bin/netca -silent -responsefile /local_dir/netca.rsp
```

In this command:

- The `-silent` option indicates to run Net Configuration Assistant in silent mode.
- `local_dir` is the full path of the directory where you copied the `netca.rsp` response file template.

A.7 Postinstallation Configuration Using Response File Created During Installation

Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.

[Using the Installation Response File for Postinstallation Configuration](#)
(page A-10)

Starting with Oracle Database 12c release 2 (12.2), you can use the response file created during installation to also complete postinstallation configuration.

[Running Postinstallation Configuration Using Response File](#) (page A-11)

Complete this procedure to run configuration assistants with the `executeConfigTools` command.

A.7.1 Using the Installation Response File for Postinstallation Configuration

Starting with Oracle Database 12c release 2 (12.2), you can use the response file created during installation to also complete postinstallation configuration.

Run the installer with the `-executeConfigTools` option to configure configuration assistants after installing Oracle Grid Infrastructure or Oracle Database. You can use the response file located at `Oracle_home/install/response/product_timestamp.rsp` to obtain the passwords required to run the configuration tools. You must update the response file with the required passwords before running the `-executeConfigTools` command.

Oracle strongly recommends that you maintain security with a password response file:

- Permissions on the response file should be set to 600.
- The owner of the response file should be the installation owner user, with the group set to the central inventory (`oraInventory`) group.

Example A-1 Response File Passwords for Oracle Grid Infrastructure

```
oracle.install.crs.config.ipmi.bmcPassword=password  
oracle.install.asm.SYSASMPassword=password  
oracle.install.asm.monitorPassword=password  
oracle.install.config.emAdminPassword=password
```

If you do not have a BMC card, or you do not want to enable IPMI, then leave the `ipmi.bmcPassword` input field blank.

If you do not want to enable Oracle Enterprise Manager for management, then leave the `emAdminPassword` password field blank.

Example A-2 Response File Passwords for Oracle Grid Infrastructure for a Standalone Server

```
oracle.install.asm.SYSASMPassword=password
oracle.install.asm.monitorPassword=password
oracle.install.config.emAdminPassword=password
```

If you do not want to enable Oracle Enterprise Manager for management, then leave the `emAdminPassword` password field blank.

Example A-3 Response File Passwords for Oracle Database

This example illustrates the passwords to specify for use with the database configuration assistants.

```
oracle.install.db.config.starterdb.password.SYS=password
oracle.install.db.config.starterdb.password.SYSTEM=password
oracle.install.db.config.starterdb.password.DBSNMP=password
oracle.install.db.config.starterdb.password.PDBADMIN=password
oracle.install.db.config.starterdb.emAdminPassword=password
oracle.install.db.config.asm.ASMSNMPPassword=password
```

You can also specify

`oracle.install.db.config.starterdb.password.ALL=password` to use the same password for all database users.

Oracle Database configuration assistants require the SYS, SYSTEM, and DBSNMP passwords for use with DBCA. You must specify the following passwords, depending on your system configuration:

- If the database uses Oracle ASM for storage, then you must specify a password for the `ASMSNMPPassword` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you create a multitenant container database (CDB) with one or more pluggable databases (PDBs), then you must specify a password for the `PDBADMIN` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.

A.7.2 Running Postinstallation Configuration Using Response File

Complete this procedure to run configuration assistants with the `executeConfigTools` command.

1. Edit the response file and specify the required passwords for your configuration. You can use the response file created during installation, located at `ORACLE_HOME/install/response/product_timestamp.rsp`. For example:

For Oracle Grid Infrastructure:

```
oracle.install.asm.SYSASMPassword=password
oracle.install.config.emAdminPassword=password
```

2. Change directory to the Oracle home containing the installation software. For example, for Oracle Grid Infrastructure:

```
cd Grid_home
```

3. Run the configuration script using the following syntax:

For Oracle Grid Infrastructure:

```
gridSetup.sh -executeConfigTools -responseFile Grid_home/install/response/  
product_timestamp.rsp
```

For Oracle Database:

```
runInstaller -executeConfigTools -responseFile ORACLE_HOME/install/response/  
product_timestamp.rsp
```

For Oracle Database, you can also run the response file located in the directory ORACLE_HOME/inventory/response/:

```
runInstaller -executeConfigTools -responseFile ORACLE_HOME/inventory/response/  
db_install.rsp
```

The postinstallation configuration tool runs the installer in the graphical user interface mode, displaying the progress of the postinstallation configuration. Specify the `[-silent]` option to run the postinstallation configuration in the silent mode.

For example, for Oracle Grid Infrastructure:

```
$ gridSetup.sh -executeConfigTools -responseFile /u01/app/12.2.0/grid/install/  
response/grid_2016-01-09_01-03-36PM.rsp [-silent]
```

For Oracle Database:

```
$ runInstaller -executeConfigTools -responseFile ORACLE_HOME/inventory/  
response/db_2016-01-09_01-03-36PM.rsp [-silent]
```

A.8 Postinstallation Configuration Using the ConfigToolAllCommands Script

You can create and run a response file configuration after installing Oracle software. The `configToolAllCommands` script requires users to create a second response file, of a different format than the one used for installing the product.

Starting with Oracle Database 12c Release 2 (12.2), the `configToolAllCommands` script is deprecated and may be desupported in a future release.

[About the Postinstallation Configuration File](#) (page A-13)

When you run a silent or response file installation, you provide information about your servers in a response file that you otherwise provide manually during a graphical user interface installation.

[Creating a Password Response File](#) (page A-14)

Review this information to create a password response file.

[Running Postinstallation Configuration Using a Password Response File](#) (page A-14)

Complete this procedure to run configuration assistants with the `configToolAllCommands` script.

Related Topics:

[Postinstallation Configuration Using Response File Created During Installation](#)
(page A-10)

Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.

A.8.1 About the Postinstallation Configuration File

When you run a silent or response file installation, you provide information about your servers in a response file that you otherwise provide manually during a graphical user interface installation.

However, the response file does not contain passwords for user accounts that configuration assistants require after software installation is complete. The configuration assistants are started with a script called `configToolAllCommands`. You can run this script in response file mode by using a password response file. The script uses the passwords to run the configuration tools in succession to complete configuration.

If you keep the password file to use for clone installations, then Oracle strongly recommends that you store the password file in a secure location. In addition, if you have to stop an installation to fix an error, then you can run the configuration assistants using `configToolAllCommands` and a password response file.

The `configToolAllCommands` password response file has the following syntax options:

- `oracle.crs` for Oracle Grid Infrastructure components or `oracle.server` for Oracle Database components that the configuration assistants configure
- `variable_name` is the name of the configuration file variable
- `value` is the desired value to use for configuration.

The command syntax is as follows:

```
internal_component_name | variable_name=value
```

For example:

```
oracle.crs | S_ASMPASSWORD=myPassWord
```

Oracle Database configuration assistants require the SYS, SYSTEM, and DBSNMP passwords for use with DBCA. You may need to specify the following additional passwords, depending on your system configuration:

- If the database is using Oracle ASM for storage, then you must specify a password for the `S_ASMSNMPPASSWORD` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you create a multitenant container database (CDB) with one or more pluggable databases (PDBs), then you must specify a password for the `S_PDBADMINPASSWORD` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.

Oracle strongly recommends that you maintain security with a password response file:

- Permissions on the response file should be set to 600.

- The owner of the response file should be the installation owner user, with the group set to the central inventory (oraInventory) group.

A.8.2 Creating a Password Response File

Review this information to create a password response file.

To create a password response file to use with the configuration assistants, perform the following steps:

1. Create a response file that has a name of the format *filename.properties*, for example:

```
$ touch pwdrsp.properties
```

2. Open the file with a text editor, and cut and paste the sample password file contents, as shown in the examples, modifying as needed.

3. Change permissions to secure the password response file. For example:

```
$ ls -al pwdrsp.properties
-rw----- 1 oracle oinstall 0 Apr 30 17:30 pwdrsp.properties
```

Example A-4 Password response file for Oracle Grid Infrastructure

```
oracle.crs|S_ASMPASSWORD=password
oracle.crs|S_OMSPASSWORD=password
oracle.crs|S_BMCPASSWORD=password
oracle.crs|S_ASMMONITORPASSWORD=password
```

If you do not have a BMC card, or you do not want to enable IPMI, then leave the S_BMCPASSWORD input field blank.

Example A-5 Password response file for Oracle Grid Infrastructure for a Standalone Server

```
oracle.crs|S_ASMPASSWORD=password
oracle.crs|S_OMSPASSWORD=password
oracle.crs|S_ASMMONITORPASSWORD=password
```

Example A-6 Password response file for Oracle Database

This example provides a template for a password response file to use with the database configuration assistants.

```
oracle.server|S_SYSPASSWORD=password
oracle.server|S_SYSTEMPASSWORD=password
oracle.server|S_EMADMINPASSWORD=password
oracle.server|S_DBSNMPPASSWORD=password
oracle.server|S_ASMSNMPPASSWORD=password
oracle.server|S_PDBADMINPASSWORD=password
```

If you do not want to enable Oracle Enterprise Manager for management, then leave those password fields blank.

A.8.3 Running Postinstallation Configuration Using a Password Response File

Complete this procedure to run configuration assistants with the configToolAllCommands script.

1. Create a password response file as described in *Creating a Password File*.

2. Change directory to `$ORACLE_HOME/cfgtoollogs`.
3. Run the configuration script using the following syntax:

```
configToolAllCommands RESPONSE_FILE=/path/name.properties
```

For example:

```
$ ./configToolAllCommands RESPONSE_FILE=/home/oracle/pwdrsp.properties
```

Completing Preinstallation Tasks Manually

Use these instructions to complete configuration tasks manually.

Oracle recommends that you use Oracle Universal Installer and Cluster Verification Utility fixup scripts to complete minimal configuration settings. If you cannot use fixup scripts, then complete minimum system settings manually.

[Configuring SSH Manually on All Cluster Nodes](#) (page B-1)

Passwordless SSH configuration is a mandatory installation requirement. SSH is used during installation to configure cluster member nodes, and SSH is used after installation by configuration assistants, Oracle Enterprise Manager, Opatch, and other features.

[Configuring Storage Device Path Persistence Using Oracle ASMLIB](#) (page B-5)

To use Oracle ASMLIB to configure Oracle ASM devices, complete the following tasks.

[Configuring Storage Device Path Persistence Manually](#) (page B-14)

You can maintain storage file path persistence by creating a rules file.

[Configuring Kernel Parameters for Linux](#) (page B-17)

These topics explain how to configure kernel parameters manually for Linux if you cannot complete them using the fixup scripts.

B.1 Configuring SSH Manually on All Cluster Nodes

Passwordless SSH configuration is a mandatory installation requirement. SSH is used during installation to configure cluster member nodes, and SSH is used after installation by configuration assistants, Oracle Enterprise Manager, Opatch, and other features.

Automatic Passwordless SSH configuration using OUI creates RSA encryption keys on all nodes of the cluster. If you have system restrictions that require you to set up SSH manually, such as using DSA keys, then use this procedure as a guide to set up passwordless SSH. If SSH is not available, then Oracle Universal Installer (OUI) attempts to use rcp instead. However, these services are disabled by default on most Linux systems.

Note: The supported version of SSH for Linux distributions is OpenSSH.

[Checking Existing SSH Configuration on the System](#) (page B-2)

To determine if SSH is running, enter the following command.

[Configuring SSH on Cluster Nodes](#) (page B-2)

You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

[Enabling SSH User Equivalency on Cluster Nodes](#) (page B-4)

After you have copied the `authorized_keys` file that contains all keys to each node in the cluster, complete the following procedure.

B.1.1 Checking Existing SSH Configuration on the System

To determine if SSH is running, enter the following command.

```
$ pgrep sshd
```

If SSH is running, then the response to this command is one or more process ID numbers. In the home directory of the installation software owner (`grid`, `oracle`), use the command `ls -al` to ensure that the `.ssh` directory is owned and writable only by the user.

You need either an RSA or a DSA key for the SSH protocol. RSA is used with the SSH 1.5 protocol, while DSA is the default for the SSH 2.0 protocol. With OpenSSH, you can use either RSA or DSA. The instructions that follow are for SSH1. If you have an SSH2 installation, and you cannot use SSH1, then refer to your SSH distribution documentation to configure SSH1 compatibility or to configure SSH2 with DSA.

B.1.2 Configuring SSH on Cluster Nodes

You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

To configure SSH, you must first create RSA or DSA keys on each cluster node, and then copy all the keys generated on all cluster node members into an authorized keys file that is identical on each node. Note that the SSH files must be readable only by `root` and by the software installation user (`oracle`, `grid`), as SSH ignores a private key file if it is accessible by others. In the examples that follow, the DSA key is used.

To configure SSH, complete the following:

[Create SSH Directory and Create SSH Keys On Each Node](#) (page B-2)

To configure SSH, you must first create RSA or DSA keys on each cluster node.

[Add All Keys to a Common `authorized_keys` File](#) (page B-3)

To configure SSH, copy all the generated keys on all cluster node members into an authorized keys file that is identical on each node.

B.1.2.1 Create SSH Directory and Create SSH Keys On Each Node

To configure SSH, you must first create RSA or DSA keys on each cluster node.

Complete the following steps on each node:

1. Log in as the software owner (in this example, the `grid` user).
2. To ensure that you are logged in as `grid`, and to verify that the user ID matches the expected user ID you have assigned to the `grid` user, enter the commands:

```
$ id
$ id grid
```

Ensure that Oracle user group and user and the user terminal window process you are using have group and user IDs are identical.

For example:

```
uid=54322(grid) gid=54321(oinstall) groups=54321(oinstall),
54322(grid,asmadmin,asmdba)
$ id grid uid=54322(grid) gid=54321(oinstall) groups=54321(oinstall),
54322(grid,asmadmin,asmdba)
```

3. If necessary, create the `.ssh` directory in the `grid` user's home directory, and set permissions on it to ensure that only the oracle user has read and write permissions:

```
$ mkdir ~/.ssh
$ chmod 700 ~/.ssh
```

Note that the SSH configuration fails if the permissions are not set to 700.

4. Enter the following command:

```
$ /usr/bin/ssh-keygen -t dsa
```

At the prompts, accept the default location for the key file (press **Enter**).

Never distribute the private key to anyone not authorized to perform Oracle software installations.

This command writes the DSA public key to the `~/.ssh/id_dsa.pub` file and the private key to the `~/.ssh/id_dsa` file.

5. Repeat steps 1 through 4 on each node that you intend to make a member of the cluster, using the DSA key.

B.1.2.2 Add All Keys to a Common `authorized_keys` File

To configure SSH, copy all the generated keys on all cluster node members into an authorized keys file that is identical on each node.

Complete the following steps:

1. On the local node, change directories to the `.ssh` directory in the Oracle Grid Infrastructure owner's home directory (typically, either `grid` or `oracle`). Then, add the DSA key to the `authorized_keys` file using the following commands:

```
$ cat id_dsa.pub >> authorized_keys
$ ls
```

In the `.ssh` directory, you should see the `id_dsa.pub` keys that you have created, and the file `authorized_keys`.

2. On the local node, use SCP (Secure Copy) or SFTP (Secure FTP) to copy the `authorized_keys` file to the `oracle` user `.ssh` directory on a remote node. The following example is with SCP, on a node called `node2`, with the Oracle Grid Infrastructure owner `grid`, where the `grid` user path is `/home/grid`:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
```

- a. You are prompted to accept a DSA key. Enter **Yes**, and you see that the node you are copying to is added to the `known_hosts` file.
- b. When prompted, provide the password for the `grid` user, which should be the same on all nodes in the cluster. The `authorized_keys` file is copied to the remote node.

Your output should be similar to the following, where xxx represents parts of a valid IP address:

```
[grid@node1 .ssh]$ scp authorized_keys node2:/home/grid/.ssh/
The authenticity of host 'node2 (xxx.xxx.173.152)' can't be established.
DSA key fingerprint is 7e:60:60:ae:40:40:d1:a6:f7:4e:zz:me:a7:48:ae:f6:7e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1,xxx.xxx.173.152' (dsa) to the list
of known hosts
grid@node2's password:
authorized_keys          100%          828           7.5MB/s       00:00
```

- Using SSH, log in to the node where you copied the `authorized_keys` file. Then change to the `.ssh` directory, and using the `cat` command, add the DSA keys for the second node to the `authorized_keys` file, clicking **Enter** when you are prompted for a password, so that passwordless SSH is set up:

```
[grid@node1 .ssh]$ ssh node2 [grid@node2 grid]$ cd .ssh [grid@node2 ssh]
$ cat id_dsa.pub >> authorized_keys
```

- Repeat steps 2 and 3 from each node to each other member node in the cluster.
- When you have added keys from each cluster node member to the `authorized_keys` file on the last node you want to have as a cluster node member, then use `scp` to copy the `authorized_keys` file with the keys from all nodes back to each cluster node member, overwriting the existing version on the other nodes. To confirm that you have all nodes in the `authorized_keys` file, enter the command `more authorized_keys`, and determine if there is a DSA key for each member node. The file lists the type of key (`ssh-dsa`), followed by the key, and then followed by the user and server. For example:

```
ssh-dsa AAAABBBB . . . = grid@node1
```

The `grid` user's `/.ssh/authorized_keys` file on every node must contain the contents from all of the `/.ssh/id_dsa.pub` files that you generated on all cluster nodes.

B.1.3 Enabling SSH User Equivalency on Cluster Nodes

After you have copied the `authorized_keys` file that contains all keys to each node in the cluster, complete the following procedure.

In this example, the Oracle Grid Infrastructure software owner is named `grid`.

Do the following:

- On the system where you want to run OUI, log in as the `grid` user.
- Use the following command syntax, where `hostname1`, `hostname2`, and so on, are the public host names (alias and fully qualified domain name) of nodes in the cluster to run SSH from the local node to each node, including from the local node to itself, and from each node to each other node:

```
[grid@nodename]$ ssh hostname1 date [grid@nodename]$ ssh hostname2
date . . .
```

At the end of this process, the public host name for each member node should be registered in the `known_hosts` file for all other cluster nodes. If you are using a remote client to connect to the local node, and you see a message similar to "Warning: No xauth data; using fake authentication data for

X11 forwarding," then this means that your authorized keys file is configured correctly, but your SSH configuration has X11 forwarding enabled. To correct this issue, see [Setting Remote Display and X11 Forwarding Configuration](#) (page 6-26).

3. Repeat step 2 on each cluster node member.

If you have configured SSH correctly, then you can now use the `ssh` or `scp` commands without being prompted for a password. For example:

```
[grid@node1 ~]$ ssh node2 date
Mon Feb 26 23:34:42 UTC 2009
[grid@node1 ~]$ ssh node1 date
Mon Feb 26 23:34:48 UTC 2009
```

If any node prompts for a password, then verify that the `~/.ssh/authorized_keys` file on that node contains the correct public keys, and that you have created an Oracle software owner with identical group membership and IDs.

B.2 Configuring Storage Device Path Persistence Using Oracle ASMLIB

To use Oracle ASMLIB to configure Oracle ASM devices, complete the following tasks.

[About Oracle ASM with Oracle ASMLIB](#) (page B-5)

Oracle ASMLIB maintains permissions and disk labels that are persistent on the storage device, so that the label is available even after an operating system upgrade.

[Installing and Configuring Oracle ASMLIB Software](#) (page B-6)

Review this information to install and configure the Oracle Automatic Storage Management library driver software manually.

[Configuring Disk Devices to Use Oracle ASMLIB](#) (page B-8)

Configure disk devices to use in an Oracle Automatic Storage Management disk group.

[Administering Oracle ASMLIB and Disks](#) (page B-10)

Review this information to administer the Oracle Automatic Storage Management library driver and disks.

[Configuring Oracle ASMLIB for Multipath Disks](#) (page B-11)

Additional configuration is required to use the Oracle Automatic Storage Management library Driver (ASMLIB) with third party vendor multipath disks.

[Deinstalling Oracle ASMLIB On Oracle Grid Infrastructure](#) (page B-14)

If Oracle ASM library driver (Oracle ASMLIB) is installed but you do not use it for device path persistence, then deinstall Oracle ASMLIB.

B.2.1 About Oracle ASM with Oracle ASMLIB

Oracle ASMLIB maintains permissions and disk labels that are persistent on the storage device, so that the label is available even after an operating system upgrade.

The Oracle Automatic Storage Management library driver simplifies the configuration and management of block disk devices by eliminating the need to rebind block disk devices used with Oracle Automatic Storage Management (Oracle ASM) each time the system is restarted.

With Oracle ASMLIB, you define the range of disks you want to have made available as Oracle ASM disks. Oracle ASMLIB maintains permissions and disk labels that are

persistent on the storage device, so that the label is available even after an operating system upgrade.

Note:

If you configure disks using Oracle ASMLIB, then you must change the disk discovery string to `ORCL:*`. If the `diskstring` is set to `ORCL:*`, or is left empty (`""`), then the installer discovers these disks.

B.2.2 Installing and Configuring Oracle ASMLIB Software

Review this information to install and configure the Oracle Automatic Storage Management library driver software manually.

Oracle ASMLIB is included with the Oracle Linux packages, and with SUSE Linux Enterprise Server. If you are a member of the Unbreakable Linux Network, then you can install the Oracle ASMLIB RPMs by subscribing to the Oracle Linux channel, and using `yum` to retrieve the most current package for your system and kernel. For additional information, see the following URL:

<http://www.oracle.com/technetwork/topics/linux/asmlib/index-101839.html>

To install and configure the Oracle Automatic Storage Management library driver software manually, perform the following steps:

1. Enter the following command to determine the kernel version and architecture of the system:

```
# uname -rm
```

2. If necessary, download the required Oracle Automatic Storage Management library driver packages from the Oracle Technology Network website:

<http://www.oracle.com/technetwork/server-storage/linux/downloads/index-088143.html>

Note:

You must install `oracleasm-support` package version 2.0.1 or later to use Oracle ASMLIB on Red Hat Enterprise Linux 5 Advanced Server. Oracle ASMLIB is already included with SUSE Linux Enterprise Server distributions.

See Also:

My Oracle Support note 1089399.1 for information about Oracle ASMLIB support with Red Hat distributions:

<https://support.oracle.com/rs?type=doc&id=1089399.1>

3. Switch to the `root` user:

```
$ su -
```

4. Install the following packages in sequence, where *version* is the version of the Oracle Automatic Storage Management library driver, *arch* is the system architecture, and *kernel* is the version of the kernel that you are using:


```
oracleasm-support-version.arch.rpm
oracleasm-kernel-version.arch.rpm
oracleasm-lib-version.arch.rpm
```

Enter a command similar to the following to install the packages:

```
# rpm -ivh oracleasm-support-version.arch.rpm \
  oracleasm-kernel-version.arch.rpm \
  oracleasm-lib-version.arch.rpm
```

For example, if you are using the Red Hat Enterprise Linux 5 AS kernel on an AMD64 system, then enter a command similar to the following:

```
# rpm -ivh oracleasm-support-2.1.3-1.el5.x86_64.rpm \
  oracleasm-2.6.18-194.26.1.el5xen-2.0.5-1.el5.x86_64.rpm \
  oracleasm-lib-2.0.4-1.el5.x86_64.rpm
```

5. Enter the following command to run the `oracleasm` initialization script with the `configure` option:

```
# /usr/sbin/oracleasm configure -i
```

Note:

The `oracleasm` command in `/usr/sbin` is the command you should use. The `/etc/init.d` path is not deprecated, but the `oracleasm` binary in that path is now used typically for internal commands.

6. Enter the following information in response to the prompts that the script displays:

Prompt	Suggested Response
Default user to own the driver interface:	<p>Standard groups and users configuration: Specify the Oracle software owner user (for example, <code>oracle</code>)</p> <p>Job role separation groups and users configuration: Specify the Grid Infrastructure software owner (for example, <code>grid</code>)</p>
Default group to own the driver interface:	<p>Standard groups and users configuration: Specify the OSDBA group for the database (for example, <code>dba</code>).</p> <p>Job role separation groups and users configuration: Specify the OSASM group for storage administration (for example, <code>asmadmin</code>).</p>
Start Oracle ASM Library driver on boot (y/n):	Enter <code>y</code> to start the Oracle Automatic Storage Management library driver when the system starts.
Scan for Oracle ASM disks on boot (y/n)	Enter <code>y</code> to scan for Oracle ASM disks when the system starts.

The script completes the following tasks:

- Creates the `/etc/sysconfig/oracleasm` configuration file
- Creates the `/dev/oracleasm` mount point
- Mounts the ASMLIB driver file system

Note:

The Oracle ASMLIB file system is not a regular file system. It is used only by the Oracle ASM library to communicate with the Oracle ASMLIB.

7. Enter the following command to load the `oracleasm` kernel module:

```
# /usr/sbin/oracleasm init
```

B.2.3 Configuring Disk Devices to Use Oracle ASMLIB

Configure disk devices to use in an Oracle Automatic Storage Management disk group.

To configure the disk devices to use in an Oracle Automatic Storage Management disk group, perform the following steps:

1. If you intend to use IDE, SCSI, or RAID devices in the Oracle Automatic Storage Management disk group, then perform the following steps:
 - a. Install or configure the disk devices that you intend to use for the disk group and restart the system.
 - b. Enter the following command to identify the device name for the disks to use:

```
# /sbin/fdisk -l
```

Depending on the type of disk, the device name can vary.

Table B-1 Device Name Formats Based on Disk Type

Disk Type	Device Name Format	Description
IDE disk	<code>/dev/hdxn</code>	In this example, <i>x</i> is a letter that identifies the IDE disk, and <i>n</i> is the partition number. For example, <code>/dev/hda</code> is the first disk on the first IDE bus.
SCSI disk	<code>/dev/sdxn</code>	In this example, <i>x</i> is a letter that identifies the IDE disk, and <i>n</i> is the partition number. For example, <code>/dev/hda</code> is the first disk on the first IDE bus.

Table B-1 (Cont.) Device Name Formats Based on Disk Type

Disk Type	Device Name Format	Description
RAID disk	/dev/rd/cxdypz /dev/ida/cxdypz	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk, and <i>z</i> is a number that identifies the partition. For example, /dev/ida/c0d1 is the second logical drive on the first controller.

To include devices in a disk group, you can specify either whole-drive device names or partition device names.

Note:

Oracle recommends that you create a single whole-disk partition on each disk to use.

- c. Use either `fdisk` or `parted` to create a single whole-disk partition on the disk devices.
2. Enter a command similar to the following to mark a disk as an Oracle Automatic Storage Management disk:

```
# /usr/sbin/oracleasm createdisk DISK1 /dev/sdb1
```

In this example, `DISK1` is a name assigned to the disk.

Note:

- The disk names you specify can contain uppercase letters, numbers, and the underscore character. They must start with an uppercase letter.
- To create a database during the installation using the Oracle Automatic Storage Management library driver, you must change the disk discovery string to `ORCL:*`.
- If you are using a multi-pathing disk driver with Oracle ASM, then make sure that you specify the correct logical device name for the disk.

3. To make the disk available on the other nodes in the cluster, enter the following command as `root` on each node:

```
# /usr/sbin/oracleasm scandisks
```

This command identifies shared disks attached to the node that are marked as Oracle ASM disks.

B.2.4 Administering Oracle ASMLIB and Disks

Review this information to administer the Oracle Automatic Storage Management library driver and disks.

To administer the Oracle Automatic Storage Management library driver and disks, use the `/usr/sbin/oracleasm` initialization script with different options, as described in the following table:

Table B-2 Disk Management Tasks Using ORACLEASM

Task	Command Example	Description
Configure or reconfigure ASMLIB	<code>oracleasm configure -i</code>	Use the <code>configure</code> option to reconfigure the Oracle Automatic Storage Management library driver, if necessary. To see command options, enter <code>oracleasm configure</code> without the <code>-i</code> flag.
Change system restart load options for ASMLIB	<code>oracleasm enable</code>	Options are <code>disable</code> and <code>enable</code> . Use the <code>disable</code> and <code>enable</code> options to change the actions of the Oracle Automatic Storage Management library driver when the system starts. The <code>enable</code> option causes the Oracle Automatic Storage Management library driver to load when the system starts.
Load or unload ASMLIB without restarting the system	<code>oracleasm restart</code>	Options are <code>start</code> , <code>stop</code> , and <code>restart</code> . Use the <code>start</code> , <code>stop</code> , and <code>restart</code> options to load or unload the Oracle Automatic Storage Management library driver without restarting the system.
Mark a disk for use with ASMLIB	<code>oracleasm createdisk VOL1 /dev/sda1</code>	Use the <code>createdisk</code> option to mark a disk device for use with the Oracle Automatic Storage Management library driver and give it a name, where <code>labelname</code> is the name you want to use to mark the device, and <code>devicepath</code> is the path to the device: <code>oracleasm createdisk labelname devicepath</code>
Unmark a named disk device	<code>oracleasm deletedisk VOL1</code>	Use the <code>deletedisk</code> option to unmark a named disk device, where <code>diskname</code> is the name of the disk: <code>oracleasm deletedisk diskname</code> Caution: Do not use this command to unmark disks that are being used by an Oracle Automatic Storage Management disk group. You must delete the disk from the Oracle Automatic Storage Management disk group before you unmark it.

Table B-2 (Cont.) Disk Management Tasks Using ORACLEASM

Task	Command Example	Description
Determine if ASMLIB is using a disk device	<code>oracleasm querydisk</code>	Use the <code>querydisk</code> option to determine if a disk device or disk name is being used by the Oracle Automatic Storage Management library driver, where <code>diskname_devicename</code> is the name of the disk or device that you want to query: <code>oracleasm querydisk diskname_devicename</code>
List Oracle ASMLIB disks	<code>oracleasm listdisks</code>	Use the <code>listdisks</code> option to list the disk names of marked Oracle ASM library driver disks.
Identify disks marked as ASMLIB disks	<code>oracleasm scandisks</code>	Use the <code>scandisks</code> option to enable cluster nodes to identify which shared disks have been marked as ASMLIB disks on another node.
Rename ASMLIB disks	<code>oracleasm renamedisk VOL1 VOL2</code>	Use the <code>renamedisk</code> option to change the label of an Oracle ASM library driver disk or device by using the following syntax, where <code>manager</code> specifies the manager device, <code>label_device</code> specifies the disk you intend to rename, as specified either by OracleASM label name or by the device path, and <code>new_label</code> specifies the new label you want to use for the disk: <code>oracleasm renamedisk [-l manager] [-v] label_device new_label</code> Use the <code>-v</code> flag to provide a verbose output for debugging. Caution: You must ensure that all Oracle Database and Oracle ASM instances have ceased using the disk before you relabel the disk. If you do not do this, then you may lose data.

B.2.5 Configuring Oracle ASMLIB for Multipath Disks

Additional configuration is required to use the Oracle Automatic Storage Management library Driver (ASMLIB) with third party vendor multipath disks.

[About Using Oracle ASM with Multipath Disks](#) (page B-12)

Oracle ASM requires that each disk is uniquely identified. If the same disk appears under multiple paths, then it causes errors.

[About Disk Scan Ordering](#) (page B-12)

The `ORACLEASM_SCANORDER` variable specifies disks to be scanned first.

[Configuring Disk Scan Ordering to Select Multipath Disks](#) (page B-13)

To configure ASMLIB to select multipath disks first, complete the following procedure.

[Configuring Disk Order Scan to Exclude Single Path Disks](#) (page B-13)

To configure ASMLIB to exclude particular single path disks, complete the following procedure.

See Also: My Oracle Support site for updates to supported storage options:

<https://support.oracle.com/>

B.2.5.1 About Using Oracle ASM with Multipath Disks

Oracle ASM requires that each disk is uniquely identified. If the same disk appears under multiple paths, then it causes errors.

In a multipath disk configuration, the same disk can appear three times: the initial path to the disk, the second path to the disk, and the multipath disk access point.

For example: If you have one local disk, `/dev/sda`, and one disk attached with external storage, then your server shows two connections, or paths, to that external storage. The Linux SCSI driver shows both paths. They appear as `/dev/sdb` and `/dev/sdc`. The system may access either `/dev/sdb` or `/dev/sdc`, but the access is to the same disk.

If you enable multipathing, then you have a multipath disk (for example, `/dev/multipatha`), which can access both `/dev/sdb` and `/dev/sdc`; any I/O to `multipatha` can use either the `sdb` or `sdc` path. If a system is using the `/dev/sdb` path, and that cable is unplugged, then the system shows an error. But the multipath disk will switch from the `/dev/sdb` path to the `/dev/sdc` path.

Most system software is unaware of multipath configurations. They can use any paths (`sdb`, `sdc` or `multipatha`). ASMLIB also is unaware of multipath configurations.

By default, ASMLIB recognizes the first disk path that Linux reports to it, but because it imprints an identity on that disk, it recognizes that disk only under one path. Depending on your storage driver, it may recognize the multipath disk, or it may recognize one of the single disk paths.

Instead of relying on the default, you should configure Oracle ASM to recognize the multipath disk.

B.2.5.2 About Disk Scan Ordering

The `ORACLEASM_SCANORDER` variable specifies disks to be scanned first.

The ASMLIB configuration file is located in the path `/etc/sysconfig/oracleasm`. It contains all the startup configuration you specified with the command `/etc/init.d/oracleasm configure`. That command cannot configure scan ordering.

The configuration file contains many configuration variables. The `ORACLEASM_SCANEXCLUDE` variable specifies the disks that are to be ignored.

Configure values for `ORACLEASM_SCANORDER` using space-delimited prefix strings. A *prefix string* is the common string associated with a type of disk. For example, if you use the prefix string `sd`, then this string matches all SCSI devices, including `/dev/sda`, `/dev/sdb`, `/dev/sdc` and so on. Note that these are not globs. They do not use wild cards. They are simple prefixes. Also note that the path is not a part of the prefix. For example, the `/dev/` path is not part of the prefix for SCSI disks that are in the path `/dev/sd*`.

For Oracle Linux and Red Hat Enterprise Linux version 5, when scanning, the kernel sees the devices as `/dev/mapper/XXX` entries. By default, the device file naming scheme `udev` creates the `/dev/mapper/XXX` names for human readability. Any configuration using `ORACLEASM_SCANORDER` should use the `/dev/mapper/XXX` entries.

B.2.5.3 Configuring Disk Scan Ordering to Select Multipath Disks

To configure ASMLIB to select multipath disks first, complete the following procedure.

1. Using a text editor, open the ASMLIB configuration file `/etc/sysconfig/oracleasm`.
2. Edit the `ORACLEASM_SCANORDER` variable to provide the prefix path of the multipath disks.

For example, if the multipath disks use the prefix `multipath` (`/dev/mapper/multipatha`, `/dev/mapper/multipathb` and so on), and the multipath disks mount SCSI disks, then provide a prefix path similar to the following:

```
ORACLEASM_SCANORDER="multipath sd"
```

3. Save the file.

When you have completed this procedure, then when ASMLIB scans disks, it first scans all disks with the prefix string `multipath`, and labels these disks as Oracle ASM disks using the `/dev/mapper/multipathX` value. It then scans all disks with the prefix string `sd`. However, because ASMLIB recognizes that these disks have already been labeled with the `/dev/mapper/multipath` string values, it ignores these disks. After scanning for the prefix strings `multipath` and `sd`, Oracle ASM then scans for any other disks that do not match the scan order.

In the example in step 2, the key word `multipath` is actually the alias for multipath devices configured in `/etc/multipath.conf` under the `multipaths` section. The default device name is in the format `/dev/mapper/mpath*` (or a similar path). For example:

```

multipaths {
    multipath {
        wwid                3600508b4000156d700012000000b0000
        alias                multipath
        ...
    }
    multipath {
        ...
        alias                mympath
        ...
    }
    ...
}

```

B.2.5.4 Configuring Disk Order Scan to Exclude Single Path Disks

To configure ASMLIB to exclude particular single path disks, complete the following procedure.

1. Using a text editor, open the ASMLIB configuration file `/etc/sysconfig/oracleasm`.
2. Edit the `ORACLEASM_SCANEXCLUDE` variable to provide the prefix path of the single path disks.

For example, if you want to exclude the single path disks `/dev/sdb` and `/dev/sdc`, then provide a prefix path similar to the following:

```
ORACLEASM_SCANEXCLUDE="sdb sdc"
```

3. Save the file.

When you have completed this procedure, then when ASMLIB scans disks, it scans all disks except for the disks with the `sdb` and `sdc` prefixes, so that it ignores `/dev/sdb` and `/dev/sdc`. It does not ignore other SCSI disks, nor multipath disks. If you have a multipath disk (for example, `/dev/multipatha`), which accesses both `/dev/sdb` and `/dev/sdc`, but you have configured ASMLIB to ignore `sdb` and `sdc`, then ASMLIB ignores these disks and instead marks only the multipath disk as an Oracle ASM disk.

B.2.6 Deinstalling Oracle ASMLIB On Oracle Grid Infrastructure

If Oracle ASM library driver (Oracle ASMLIB) is installed but you do not use it for device path persistence, then deinstall Oracle ASMLIB.

1. Log in as `root`.
2. Stop Oracle ASM and any running database instance on the node:

```
$ srvctl stop instance -d db_unique_name-node node_name
$ srvctl stop asm -node node_name
```

To stop the last Oracle Flex ASM instance on the node, stop the Oracle Clusterware stack:

```
Grid_home/bin/crsctl stop crs
```

3. Stop Oracle ASMLIB:

```
/etc/init.d/oracleasm disable
```

4. Remove the `oracleasm` library and tools RPMs:

```
# rpm -e oracleasm-support
# rpm -e oracleasm
```

5. Check if any `oracleasm` RPMs are remaining:

```
# rpm -qa | grep oracleasm
```

6. If any `oracleasm` configuration files are remaining, remove them:

```
# rpm -qa | grep oracleasm | xargs rpm -e
```

Oracle ASMLIB and associated RPMs are now removed.

Start the Oracle Clusterware stack. Optionally, you can install and configure Oracle ASM Filter Driver (Oracle ASMTFD) before starting the Oracle Clusterware stack.

Oracle Automatic Storage Management Administrator's Guide

B.3 Configuring Storage Device Path Persistence Manually

You can maintain storage file path persistence by creating a rules file.

By default, the device file naming scheme `udev` dynamically creates device file names when the server is started, and assigns ownership of them to `root`. If `udev` applies default settings, then it changes Oracle device file names and owners for the disks, making the disks inaccessible when the server is restarted. For example, a voting disk

on a device named `/dev/sdd` owned by the user `grid` may be on a device named `/dev/sdf` owned by `root` after restarting the server.

If you use Oracle ASMFD, then you do not have to ensure permissions and device path persistence in `udev`.

If you do not use Oracle ASMFD, then you must create a custom rules file. Linux vendors customize their `udev` configurations and use different orders for reading rules files. For example, on some Linux distributions when `udev` is started, it sequentially carries out rules (configuration directives) defined in rules files. These files are in the path `/etc/udev/rules.d/`. Rules files are read in lexical order. For example, rules in the file `10-wacom.rules` are parsed and carried out before rules in the rules file `90-ib.rules`.

When specifying the device information in the `udev` rules file, ensure that the `OWNER`, `GROUP`, and `MODE` are specified before all other characteristics in the order shown. For example, to include the characteristic `ACTION` on the `UDEV` line, specify `ACTION` after `OWNER`, `GROUP`, and `MODE`.

Where rules files describe the same devices, on the supported Linux kernel versions, the last file read is the one that is applied.

[Configuring Device Persistence Manually for Oracle ASM](#) (page B-15)

Complete these tasks to create device path persistence manually for Oracle ASM.

B.3.1 Configuring Device Persistence Manually for Oracle ASM

Complete these tasks to create device path persistence manually for Oracle ASM.

1. On existing systems (single instance or clustered), to obtain current mount paths, run the command `scsi_id (/sbin/scsi_id)` on storage devices on a server to obtain their unique device identifiers. When you run the command `scsi_id` with the `-s` argument, the device path and name you provide should be relative to the `sysfs` directory `/sys` (for example, `/block/device`) when referring to `/sys/block/device`.

For example:

```
# /sbin/scsi_id -g -s /block/sdb/sdb1
360a98000686f6959684a453333524174
```

```
# /sbin/scsi_id -g -s /block/sde/sde1
360a98000686f6959684a453333524179
```

Record the unique SCSI identifiers, so you can provide them when required.

Note:

The command `scsi_id` should return the same device identifier value for a given device, regardless of which node the command is run from.

2. Configure SCSI devices as trusted devices (white listed), by editing the `/etc/scsi_id.config` file and adding `options=-g` to the file. For example:

```
# cat > /etc/scsi_id.config
vendor="ATA",options=-p 0x80
options=-g
```

- Using a text editor, create a UDEV rules file for the Oracle ASM devices, setting permissions to 0660 for the installation owner and the operating system group you have designated the OSASM group, whose members are administrators of the Oracle Grid Infrastructure software. For example, on Oracle Linux, to create a role-based configuration rules .d file where the installation owner is grid and the OSASM group asmadmin, enter commands similar to the following:

```
# vi /etc/udev/rules.d/99-oracle-asmdevices.rules

KERNEL=="sdb1", OWNER="grid", GROUP="asmadmin", MODE="0660",
BUS=="scsi", PROGRAM==" /sbin/scsi_id", RESULT=="14f70656e66696c00000000"
KERNEL=="sdcl", OWNER="grid", GROUP="asmadmin", MODE="0660",
BUS=="scsi", PROGRAM==" /sbin/scsi_id", RESULT=="14f70656e66696c00000001"
KERNEL=="sdd1", OWNER="grid", GROUP="asmadmin", MODE="0660",
BUS=="scsi", PROGRAM==" /sbin/scsi_id", RESULT=="14f70656e66696c00000002"
```

- On clustered systems, copy the rules .d file to all other nodes on the cluster. For example:

```
# scp 99-oracle-asmdevices.rules root@node2:/etc/udev/rules.d/99-oracle-
asmdevices.rules
```

- As root, run the partprobe command using the syntax /sbin/partprobe *devicename*.

For example:

```
# /sbin/partprobe /dev/sdcl
# /sbin/partprobe /dev/sdd1
# /sbin/partprobe /dev/sdel
# /sbin/partprobe /dev/sdf1
```

- Run the command udevtest (/sbin/udevtest) to test the UDEV rules configuration you have created. The output should indicate that the devices are available and the rules are applied as expected. For example, for /dev/ssd1:

```
# udevtest /block/sdd/sdd1
main: looking at device '/block/sdd/sdd1' from subsystem 'block'
udev_rules_get_name: add symlink
'disk/by-id/scsi-360a98000686f6959684a453333524174-part1'
udev_rules_get_name: add symlink
'disk/by-path/ip-192.168.1.1:3260-iscsi-iqn.1992-08.com.netapp:sn.887085-part1'
udev_node_mknod: preserve file '/dev/.tmp-8-17', because it has correct dev_t
run_program: '/lib/udev/vol_id --export /dev/.tmp-8-17'
run_program: '/lib/udev/vol_id' returned with status 4
run_program: '/sbin/scsi_id'
run_program: '/sbin/scsi_id' (stdout) '360a98000686f6959684a453333524174'
run_program: '/sbin/scsi_id' returned with status 0
udev_rules_get_name: rule applied, 'sdd1' becomes 'data1'
udev_device_event: device '/block/sdd/sdd1' validate currently present symlinks
udev_node_add: creating device node '/dev/data1', major = '8', minor = '17',
mode = '0640', uid = '0', gid = '500'
udev_node_add: creating symlink
'/dev/disk/by-id/scsi-360a98000686f6959684a453333524174-part1' to '../data1'
udev_node_add: creating symlink
'/dev/disk/by-path/ip-192.168.1.1:3260-iscsi-iqn.1992-08.com.netapp:sn.84187085
-part1' to '../data1'
main: run: 'socket:/org/kernel/udev/monitor'
main: run: '/lib/udev/udev_run_devd'
main: run: 'socket:/org/freedesktop/hal/udev_event'
main: run: '/sbin/pam_console_apply /dev/data1
/dev/disk/by-id/scsi-360a98000686f6959684a453333524174-part1'
```

```
/dev/disk/by-path/ip-192.168.1.1:3260-iscsi-iqn.1992-08.com.netapp:sn.84187085-
part1'
```

In the example output, note that applying the rules renames OCR device `/dev/sdd1` to `/dev/data1`.

7. Load the rules and restart the UDEV service. For example:

- Oracle Linux and Red Hat Enterprise Linux


```
# udevadm control --reload-rules
```
- SUSE Linux Enterprise Server


```
# /etc/init.d boot.udev restart
```

Verify that the device permissions and ownerships are set correctly.

B.4 Configuring Kernel Parameters for Linux

These topics explain how to configure kernel parameters manually for Linux if you cannot complete them using the fixup scripts.

[Minimum Parameter Settings for Installation](#) (page B-17)

Use this table to set parameters manually if you cannot use the fixup scripts.

[Changing Kernel Parameter Values](#) (page B-19)

Use these instructions to display and change the kernel parameter values if they are different from the minimum recommended value.

[Configuring Additional Kernel Settings for SUSE Linux](#) (page B-20)

Additional kernel parameter settings configuration is required on SUSE Linux Enterprise Server systems.

[Setting UDP and TCP Kernel Parameters Manually](#) (page B-21)

If you do not use a Fixup script or CVU to set ephemeral ports, then set TCP/IP ephemeral port range parameters to provide enough ephemeral ports for the anticipated server workload.

B.4.1 Minimum Parameter Settings for Installation

Use this table to set parameters manually if you cannot use the fixup scripts.

Note:

- Unless otherwise specified, the kernel parameter and shell limit values shown in the following table are minimum values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See the operating system documentation for more information about tuning kernel parameters.
 - If the current value for any parameter is greater than the value listed in this table, then the Fixup scripts do not change the value of that parameter.
-
-

Table B-3 Minimum Operating System Resource Parameter Settings

Parameter	Value	File
semmsl	250	/proc/sys/ kernel/sem
semmns	32000	
semopm	100	
semmni	128	
shmall	40 percent of the size of physical memory in pages Note: If the server supports multiple databases, or uses a large SGA, then set this parameter to a value that is equal to the total amount of shared memory, in 4K pages, that the system can use at one time.	/proc/sys/kernel/ shmall
shmmax	Half the size of physical memory in bytes See My Oracle Support Note 567506.1 for additional information about configuring shmmax.	/proc/sys/kernel/ shmmax
shmmni	4096	/proc/sys/kernel/ shmmni
panic_on_oops	1	/proc/sys/kernel/ panic_on_oops
file-max	6815744	/proc/sys/fs/file- max
aio-max-nr	1048576 Note: This value limits concurrent outstanding requests and should be set to avoid I/O subsystem failures.	/proc/sys/fs/aio- max-nr
ip_local_port_range	Minimum: 9000 Maximum: 65500	/proc/sys/net/ipv4/ ip_local_port_range
rmem_default	262144	/proc/sys/net/core/ rmem_default
rmem_max	4194304	/proc/sys/net/core/ rmem_max
wmem_default	262144	/proc/sys/net/core/ wmem_default
wmem_max	1048576	/proc/sys/net/core/ wmem_max

B.4.2 Changing Kernel Parameter Values

Use these instructions to display and change the kernel parameter values if they are different from the minimum recommended value.

Use the following commands to display the current values of the kernel parameters:

Table B-4 *Commands to Display Kernel Parameter Values*

Parameter	Command
semmsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm This command displays the details of the shared memory segment sizes.
file-max	# /sbin/sysctl -a grep file-max This command displays the maximum number of file handles.
ip_local_port_range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.
rmem_default	# /sbin/sysctl -a grep rmem_default
rmem_max	# /sbin/sysctl -a grep rmem_max
wmem_default	# /sbin/sysctl -a grep wmem_default
wmem_max	# /sbin/sysctl -a grep wmem_max
aio-max-nr	# /sbin/sysctl -a grep aio-max-nr

If you used the Oracle Preinstallation RPM to complete your preinstallation configuration tasks, then the Oracle Preinstallation RPM sets these kernel parameters for you. However, if you did not use the Oracle Preinstallation RPM or the kernel parameters are different from the minimum recommended value, then to change these kernel parameter values:

1. Using any text editor, create or edit the `/etc/sysctl.d/97-oracle-database-sysctl.conf` file, and add or edit lines similar to:

```
fs.aio-max-nr = 1048576
fs.file-max = 6815744
kernel.shmall = 2097152
kernel.shmmax = 4294967295
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 262144
```

```
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
```

Note:

- If you used the Oracle Preinstallation RPM, then your kernel parameter settings reside in the `/etc/sysctl.d/99-oracle-database-server-12cR2-preinstall-sysctl.conf` file.
 - Include lines only for the kernel parameter values to change. For the semaphore parameters (`kernel.sem`), you must specify all four values. If any of the current values are larger than the minimum value, then specify the larger value.
 - The `/etc/sysctl.conf` file has been deprecated.
-

By specifying the values in the `/etc/sysctl.d/97-oracle-database-sysctl.conf` file, the values persist on system restarts.

2. To change the current values of the kernel parameters:

```
# /sbin/sysctl --system
```

Review the output. If the values are incorrect, edit the `/etc/sysctl.d/97-oracle-database-sysctl.conf` file, then enter this command again.

3. Confirm that the values are set correctly:

```
# /sbin/sysctl -a
```

4. Restart the computer, or run `sysctl --system` to make the changes in the `/etc/sysctl.d/97-oracle-database-sysctl.conf` file available in the active kernel memory.

B.4.3 Configuring Additional Kernel Settings for SUSE Linux

Additional kernel parameter settings configuration is required on SUSE Linux Enterprise Server systems.

Perform these steps on SUSE Linux Enterprise Server systems only:

1. Enter the following command to enable the system to read the `/etc/sysctl.conf` file when it restarts:

```
# /sbin/chkconfig boot.sysctl on
```

2. Enter the GID of the `oinstall` group as the value for the parameter `/proc/sys/vm/hugetlb_shm_group`.

For example, where the `oinstall` group GID is 501:

```
# echo 501 > /proc/sys/vm/hugetlb_shm_group
```

Use `vi` to add the following text to `/etc/sysctl.conf`, and enable the `boot.sysctl` script to run on system restart:

```
vm.hugetlb_shm_group=501
```

Note: Only one group can be defined as the `vm.hugetlb_shm_group`.

B.4.4 Setting UDP and TCP Kernel Parameters Manually

If you do not use a Fixup script or CVU to set ephemeral ports, then set TCP/IP ephemeral port range parameters to provide enough ephemeral ports for the anticipated server workload.

Ensure that the lower range is set to at least 9000 or higher, to avoid Well Known ports, and to avoid ports in the Registered Ports range commonly used by Oracle and other server ports. Set the port range high enough to avoid reserved ports for any applications you may intend to use. If the lower value of the range you have is greater than 9000, and the range is large enough for your anticipated workload, then you can ignore Oracle Universal Installer warnings regarding the ephemeral port range.

For example, with IPv4, use the following command to check your current range for ephemeral ports:

```
$ cat /proc/sys/net/ipv4/ip_local_port_range
32768 61000
```

In the preceding example, the lowest port (32768) and the highest port (61000) are set to the default range.

If necessary, update the UDP and TCP ephemeral port range to a range high enough for anticipated system workloads, and to ensure that the ephemeral port range starts at 9000 and above. For example:

```
# echo 9000 65500 > /proc/sys/net/ipv4/ip_local_port_range
```

Oracle recommends that you make these settings permanent. For example, as `root`, use a text editor to open `/etc/sysctl.conf`, and add or change to the following:
`net.ipv4.ip_local_port_range = 9000 65500`, and then restart the network:

```
# /etc/rc.d/init.d/network restart
```

Refer to your Linux distribution system administration documentation for information about automating ephemeral port range alteration on system restarts.

Optimal Flexible Architecture

Oracle Optimal Flexible Architecture (OFA) rules are a set of configuration guidelines created to ensure well-organized Oracle installations, which simplifies administration, support and maintenance.

[About the Optimal Flexible Architecture Standard](#) (page C-1)

Oracle Optimal Flexible Architecture (OFA) rules help you to organize database software and configure databases to allow multiple databases, of different versions, owned by different users to coexist.

[About Multiple Oracle Homes Support](#) (page C-2)

Oracle Database supports multiple Oracle homes. You can install this release or earlier releases of the software more than once on the same system, in different Oracle home directories.

[About the Oracle Inventory Directory and Installation](#) (page C-3)

The directory that you designate as the Oracle Inventory directory (oraInventory) stores an inventory of all software installed on the system.

[Oracle Base Directory Naming Convention](#) (page C-4)

This section describes what the Oracle base is, and how it functions.

[Oracle Home Directory Naming Convention](#) (page C-4)

By default, Oracle Universal Installer configures Oracle home directories using these Oracle Optimal Flexible Architecture conventions.

[Optimal Flexible Architecture File Path Examples](#) (page C-5)

This topic shows examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.

C.1 About the Optimal Flexible Architecture Standard

Oracle Optimal Flexible Architecture (OFA) rules help you to organize database software and configure databases to allow multiple databases, of different versions, owned by different users to coexist.

In earlier Oracle Database releases, the OFA rules provided optimal system performance by isolating fragmentation and minimizing contention. In current releases, OFA rules provide consistency in database management and support, and simplifies expanding or adding databases, or adding additional hardware.

By default, Oracle Universal Installer places Oracle Database components in directory locations and with permissions in compliance with OFA rules. Oracle recommends that you configure all Oracle components on the installation media in accordance with OFA guidelines.

Oracle recommends that you accept the OFA default. Following OFA rules is especially of value if the database is large, or if you plan to have multiple databases.

Note:

OFA assists in identification of an ORACLE_BASE with its Automatic Diagnostic Repository (ADR) diagnostic data to properly collect incidents.

C.2 About Multiple Oracle Homes Support

Oracle Database supports multiple Oracle homes. You can install this release or earlier releases of the software more than once on the same system, in different Oracle home directories.

Careful selection of mount point names can make Oracle software easier to administer. Configuring multiple Oracle homes in compliance with Optimal Flexible Architecture (OFA) rules provides the following advantages:

- You can install this release, or earlier releases of the software, more than once on the same system, in different Oracle home directories. However, you cannot install products from one release of Oracle Database into an Oracle home directory of a different release. For example, you cannot install Oracle Database 12c software into an existing Oracle 11g Oracle home directory.
- Multiple databases, of different versions, owned by different users can coexist concurrently.
- You must install a new Oracle Database release in a new Oracle home that is separate from earlier releases of Oracle Database.

You cannot install multiple releases in one Oracle home. Oracle recommends that you create a separate Oracle Database Oracle home for each release, in accordance with the Optimal Flexible Architecture (OFA) guidelines.

- In production, the Oracle Database server software release must be the same as the Oracle Database dictionary release through the first four digits (the major, maintenance, and patch release number).
- Later Oracle Database releases can access earlier Oracle Database releases. However, this access is only for upgrades. For example, Oracle Database 12c release 2 can access an Oracle Database 11g release 2 (11.2.0.4) database if the 11.2.0.4 database is started up in upgrade mode.
- Oracle Database Client can be installed in the same Oracle Database home if both products are at the same release level. For example, you can install Oracle Database Client 12.2.0.1 into an existing Oracle Database 12.2.0.1 home but you cannot install Oracle Database Client 12.2.0.1 into an existing Oracle Database 12.1.0.2 home. If you apply a patch set before installing the client, then you must apply the patch set again.
- Structured organization of directories and files, and consistent naming for database files simplify database administration.
- Login home directories are not at risk when database administrators add, move, or delete Oracle home directories.
- You can test software upgrades in an Oracle home in a separate directory from the Oracle home where your production database is located.

C.3 About the Oracle Inventory Directory and Installation

The directory that you designate as the Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system.

All Oracle software installation owners on a server are granted the `OINSTALL` privileges to read and write to this directory. If you have previous Oracle software installations on a server, then additional Oracle software installations detect this directory from the `/etc/oraInst.loc` file, and continue to use that Oracle Inventory. Ensure that the group designated as the `OINSTALL` group is available as a primary group for all planned Oracle software installation owners.

If you are installing Oracle software for the first time, then OUI creates an Oracle base and central inventory, and creates an Oracle inventory using information in the following priority:

- In the path indicated in the `ORACLE_BASE` environment variable set for the installation owner user account
- In an Optimal Flexible Architecture (OFA) path (`u[01-99]/app/owner` where *owner* is the name of the user account running the installation), and that user account has permissions to write to that path
- In the user home directory, in the path `/app/owner`, where *owner* is the name of the user account running the installation

For example:

If you are performing an Oracle Database installation, and you set `ORACLE_BASE` for user `oracle` to the path `/u01/app/oracle` before installation, and grant 755 permissions to `oracle` for that path, then Oracle Universal Installer creates the Oracle Inventory directory one level above the `ORACLE_BASE` in the path `ORACLE_BASE/.. /oraInventory`, so the Oracle Inventory path is `/u01/app/oraInventory`. Oracle Universal Installer installs the software in the `ORACLE_BASE` path. If you are performing an Oracle Grid Infrastructure for a Cluster installation, then the Grid installation path is changed to `root` ownership after installation, and the Grid home software location should be in a different path from the Grid user Oracle base.

If you create the OFA path `/u01`, and grant `oracle` 755 permissions to write to that path, then the Oracle Inventory directory is created in the path `/u01/app/oraInventory`, and Oracle Universal Installer creates the path `/u01/app/oracle`, and configures the `ORACLE_BASE` environment variable for the Oracle user to that path. If you are performing an Oracle Database installation, then the Oracle home is installed under the Oracle base. However, if you are installing Oracle Grid Infrastructure for a cluster, then be aware that ownership of the path for the Grid home is changed to `root` after installation and the Grid base and Grid home should be in different locations, such as `/u01/grid` for the Grid home path, and `/u01/app/grid` for the Grid base. For example:

```
/u01/app/oraInventory, owned by grid:oinstall
/u01/app/oracle, owned by oracle:oinstall
/u01/app/oracle/product/12.2.0/dbhome_1/, owned by
oracle:oinstall
/u01/app/grid, owned by grid:oinstall
/u01/app/12.2.0/grid, owned by root
```

If you have neither set `ORACLE_BASE`, nor created an OFA-compliant path, then the Oracle Inventory directory is placed in the home directory of the user that is performing the installation, and the Oracle software is installed in the path `/app/owner`, where *owner* is the Oracle software installation owner. For example:

```
/home/oracle/oraInventory
/home/oracle/app/oracle/product/12.2.0/dbhome_1
```

C.4 Oracle Base Directory Naming Convention

This section describes what the Oracle base is, and how it functions.

The Oracle Base directory is the database home directory for Oracle Database installation owners, and the log file location for Oracle Grid Infrastructure owners. Name Oracle base directories using the syntax `/pm/h/u`, where *pm* is a string mount point name, *h* is selected from a small set of standard directory names, and *u* is the name of the owner of the directory.

You can use the same Oracle base directory for multiple installations. If different operating system users install Oracle software on the same system, then you must create a separate Oracle base directory for each installation owner. For ease of administration, Oracle recommends that you create a unique owner for each Oracle software installation owner, to separate log files.

Because all Oracle installation owners write to the central Oracle inventory file, and that file mountpoint is in the same mount point path as the initial Oracle installation, Oracle recommends that you use the same `/pm/h` path for all Oracle installation owners.

Table C-1 Examples of OFA-Compliant Oracle Base Directory Names

Example	Description
<code>/u01/app/oracle</code>	Oracle Database Oracle base, where the Oracle Database software installation owner name is <code>oracle</code> . The Oracle Database binary home is located underneath the Oracle base path.
<code>/u01/app/grid</code>	Oracle Grid Infrastructure Oracle base, where the Oracle Grid Infrastructure software installation owner name is <code>grid</code> .
	Caution: The Oracle Grid Infrastructure Oracle base should not contain the Oracle Grid Infrastructure binaries for an Oracle Grid Infrastructure for a cluster installation. Permissions for the file path to the Oracle Grid Infrastructure binary home is changed to <code>root</code> during installation.

C.5 Oracle Home Directory Naming Convention

By default, Oracle Universal Installer configures Oracle home directories using these Oracle Optimal Flexible Architecture conventions.

The directory pattern syntax for Oracle homes is `/pm/s/u/product/v/type_[n]`. The following table describes the variables used in this syntax:

Variable	Description
<i>pm</i>	A mount point name.

Variable	Description
<i>s</i>	A standard directory name.
<i>u</i>	The name of the owner of the directory.
<i>v</i>	The version of the software.
<i>type</i>	The type of installation. For example: Database (<i>dbhome</i>), Client (<i>client</i>), or Oracle Grid Infrastructure (<i>grid</i>)
<i>n</i>	An optional counter, which enables you to install the same product more than once in the same Oracle base directory. For example: Database 1 and Database 2 (<i>dbhome_1</i> , <i>dbhome_2</i>)

For example, the following path is typical for the first installation of Oracle Database on this system:

```
/u01/app/oracle/product/12.2.0/dbhome_1
```

C.6 Optimal Flexible Architecture File Path Examples

This topic shows examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.

This example shows an Optimal Flexible Architecture-compliant installation with three Oracle home directories and three databases, as well as examples of the deployment path differences between a cluster install and a standalone server install of Oracle Grid Infrastructure. The database files are distributed across three mount points: /u02, /u03, and /u04.

Note:

- The Grid homes are examples of Grid homes used for an Oracle Grid Infrastructure for a standalone server deployment (Oracle Restart), or a Grid home used for an Oracle Grid Infrastructure for a cluster deployment (Oracle Clusterware). You can have either an Oracle Restart deployment, or an Oracle Clusterware deployment. You cannot have both options deployed at the same time.
 - Oracle Automatic Storage Management (Oracle ASM) is included as part of an Oracle Grid Infrastructure installation. Oracle recommends that you use Oracle ASM to provide greater redundancy and throughput.
-
-

Table C-2 Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/	Root directory
/u01/	User data mount point 1

Table C-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/u01/app/	Subtree for application software
/u01/app/oraInventory	Central OraInventory directory, which maintains information about Oracle installations on a server. Members of the group designated as the OINSTALL group have permissions to write to the central inventory. All Oracle software installation owners must have the OINSTALL group as their primary group, and be able to write to this group.
/u01/app/oracle/	Oracle base directory for user <code>oracle</code> . There can be many Oracle Database installations on a server, and many Oracle Database software installation owners. Oracle software homes that an Oracle installation owner owns should be located in the Oracle base directory for the Oracle software installation owner, unless that Oracle software is Oracle Grid Infrastructure deployed for a cluster.
/u01/app/grid	Oracle base directory for user <code>grid</code> . The Oracle home (Grid home) for Oracle Grid Infrastructure for a cluster installation is located outside of the Grid user. There can be only one Grid home on a server, and only one Grid software installation owner. The Grid home contains log files and other administrative files.
/u01/app/oracle/admin/	Subtree for database administration files
/u01/app/oracle/admin/TAR	Subtree for support log files
/u01/app/oracle/admin/db_sales/	admin subtree for database named "sales"
/u01/app/oracle/admin/db_dwh/	admin subtree for database named "dwh"
/u01/app/oracle/fast_recovery_area/	Subtree for recovery files
/u01/app/oracle/fast_recovery_area/db_sales	Recovery files for database named "sales"

Table C-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
/u01/app/oracle/ fast_recovery_area/ db_dwh	Recovery files for database named "dwh"
/u02/app/oracle/ oradata /u03/app/oracle/ oradata /u04/app/oracle/ oradata	Oracle data file directories
/u01/app/oracle/ product/	Common path for Oracle software products other than Oracle Grid Infrastructure for a cluster
/u01/app/oracle/ product/12.2.0/ dbhome_1	Oracle home directory for Oracle Database 1, owned by Oracle Database installation owner account <code>oracle</code>
/u01/app/oracle/ product/12.2.0/ dbhome_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account <code>oracle</code>
/u01/app/oracle2/ product/12.2.0/ dbhome_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account <code>oracle2</code>
/u01/app/oracle/ product/12.2.0/grid	Oracle home directory for Oracle Grid Infrastructure for a standalone server, owned by Oracle Database and Oracle Grid Infrastructure installation owner <code>oracle</code> .
/u01/app/12.2.0/ grid	Oracle home directory for Oracle Grid Infrastructure for a cluster (Grid home), owned by user <code>grid</code> before installation, and owned by <code>root</code> after installation.

A

- adding Oracle ASM listener, [11-20](#)
- ASM_DISKSTRING, [8-12](#)
- ASMCA
 - Used to create disk groups for older Oracle Database releases on Oracle ASM, [10-10](#)
- asmdba groups
 - creating, [6-12](#)
- asmoper group
 - creating, [6-13](#)
- ASMSNMP, [1-5](#)
- Automatic Diagnostic Repository (ADR), [C-1](#)
- Automatic Storage Management Cluster File System
 - See Oracle ACFS.

B

- backupdba group
 - creating, [6-14](#)
- Bash shell
 - default user startup file, [6-23](#)
- bash_profile file, [6-23](#)
- batch upgrade, [11-12](#)
- binaries
 - relinking, [10-10](#)
- binary files
 - supported storage options for, [7-2](#)
- BMC
 - configuring, [6-30](#)
- BMC interface
 - preinstallation tasks, [6-28](#)
- Bourne shell
 - default user startup file, [6-23](#)

C

- C shell
 - default user startup file, [6-23](#)
- central inventory, [C-5](#)
 - See also Oracle inventory directory
 - See also OINSTALL directory
- changing kernel parameter values, [B-19](#)
- checkdir error, [10-10](#), [11-4](#)

- checklists, [1-1](#)
- client-server configurations, [C-2](#)
- clients
 - and upgrades, [5-5](#)
 - connecting to SCAN, [5-5](#)
 - using SCAN, [5-5](#)
- cluster configuration
 - Oracle Domain Services Cluster, [9-3](#)
 - Oracle Extended Clusters, [9-6](#)
 - Oracle Member Clusters, [9-5](#)
 - Oracle Standalone Clusters, [9-3](#)
- cluster file system
 - storage option for data files, [7-6](#)
- cluster name
 - requirements for, [1-5](#)
- cluster nodes
 - private network node interfaces, [1-5](#)
 - private node names, [5-5](#)
 - public network node names and addresses, [1-5](#)
 - virtual node names, [1-5](#), [5-4](#)
- Cluster Time Synchronization Service, [4-32](#)
- CLUSTER_INTERCONNECTS parameter, [5-3](#)
- clusterware
 - requirements for third party clusterware, [1-5](#)
- commands
 - asmca, [8-22](#)
 - asmcmd, [8-10](#)
 - crsctl, [11-4](#)
 - df -h, [2-2](#)
 - free, [2-2](#)
 - grep MemTotal, [2-2](#)
 - grep SwapTotal, [2-2](#)
 - gridSetup.sh, [9-28](#)
 - ipmitool, [6-31](#)
 - lsmod, [6-29](#)
 - modprobe, [6-29](#)
 - nscd, [4-30](#)
 - root.sh, [10-3](#)
 - rootcrs.pl
 - and deconfig option, [12-13](#)
 - rootcrs.sh, [10-10](#)
 - rootupgrade.sh, [11-4](#)
 - runcluvfy.sh, [9-28](#)

commands (*continued*)
 srvctl, [11-4](#)
 sysctl, [8-17](#)
 umask, [6-22](#)
 uname —m, [2-2](#)
 unset, [11-8](#)
 useradd, [6-15](#)
cron jobs, [1-12](#)
ctsdd, [4-32](#)
custom database
 failure groups for Oracle ASM, [8-3](#)
 requirements when using Oracle ASM, [8-3](#)

D

data files
 storage options, [7-6](#)
 supported storage options for, [7-2](#)
data loss
 minimizing with Oracle ASM, [8-3](#), [8-15](#)
Database Configuration Assistant
 running in silent mode, [A-7](#)
databases
 Oracle ASM requirements, [8-3](#)
DB_RECOVERY_FILE_DEST, [10-4](#)
DB_RECOVERY_FILE_DEST_SIZE, [10-4](#)
dba group
 creating, [6-13](#)
 description, [6-9](#)
 SYSDBA privilege, [6-9](#)
dba groups
 creating, [6-14](#)
DBCA
 no longer used for Oracle ASM disk group
 administration, [10-10](#)
 Used to create server pools for earlier Oracle
 Database releases, [10-9](#)
dbca.rsp file, [A-4](#)
default file mode creation mask
 setting, [6-22](#)
default Linux installation
 recommendation for, [4-4](#)
deinstall
 Oracle Member Cluster, [12-14](#)
 See also removing Oracle software
deinstallation
 examples, [12-6](#)
deinstallation tool, [12-2](#)
Deinstallation tool
 Restriction for Oracle Flex Clusters and -lastnode
 flag, [12-13](#)
device names
 IDE disks, [B-8](#)
 RAID, [B-8](#)
 SCSI disks, [B-8](#)
df command, [6-23](#)
dgdba group

dgdba group (*continued*)
 creating, [6-14](#)
DHCP
 and GNS, [5-10](#)
diagnostic data, [C-1](#)
Direct NFS
 disabling, [8-20](#)
 enabling, [8-20](#)
 oranfstab file, [8-18](#)
directory
 creating separate data file directories, [8-15](#)
disk group
 Oracle ASM, [8-3](#)
 recommendations for Oracle ASM disk groups,
 [8-3](#)
disk group corruption
 preventing, [8-14](#)
disk groups
 checking, [8-10](#)
 recommendations for, [8-3](#)
disk space
 requirements for preconfigured database in
 Oracle ASM, [8-3](#)
disks
 checking availability for Oracle Automatic Storage
 Management, [B-8](#)
 displaying attached disks, [B-8](#)
 selecting for use with Oracle ASM, [8-11](#)
display variable, [1-7](#)
downgrade, [11-23](#)
downgrade after failed installation, [11-27](#), [11-28](#)
downgrade after failed upgrade, [11-27](#), [11-28](#)
downgrades, [11-24](#), [11-27](#), [11-28](#)
downgrades restrictions, [11-24](#)
downgrading
 Oracle Grid Infrastructure, [11-24](#)
 to 12.1, [11-24](#)

E

enterprise.rsp file, [A-4](#)
environment
 configuring for Oracle user, [6-22](#)
environment variables
 ORACLE_BASE, [6-23](#)
 ORACLE_HOME, [6-23](#), [11-8](#)
 ORACLE_SID, [6-23](#), [11-8](#)
 removing from shell startup file, [6-23](#)
 SHELL, [6-23](#)
 TEMP and TMPDIR, [6-23](#)
errors
 X11 forwarding, [6-26](#), [B-4](#)
errors using Opatch, [11-4](#)
errors using OPatch, [10-10](#)
Exadata
 relinking binaries example for, [10-10](#)
examples

examples (*continued*)
Oracle ASM failure groups, [8-3](#)
executeConfigTools, [A-11](#)

F

failed install, [11-29](#)
failed upgrade, [11-29](#)
failure group
 characteristics of Oracle ASM failure group, [8-3](#),
 [8-15](#)
 examples of Oracle ASM failure groups, [8-3](#)
 Oracle ASM, [8-3](#)
fast recovery area
 filepath, [C-5](#)
 Grid home
 filepath, [C-5](#)
fdisk command, [B-8](#)
fencing
 and IPMI, [6-28](#)
file mode creation mask
 setting, [6-22](#)
file system
 storage option for data files, [7-6](#)
files
 bash_profile, [6-23](#)
 dbca.rsp, [A-4](#)
 editing shell startup file, [6-23](#)
 enterprise.rsp, [A-4](#)
 login, [6-23](#)
 profile, [6-23](#)
 response files, [A-3](#)
filesets, [4-10](#)

G

GIMR, [8-10](#)
globalization, [1-12](#)
GNS
 about, [5-11](#)
 configuration example, [5-22](#)
 configuring, [5-10](#)
GNS client clusters
 and GNS client data file, [5-13](#)
 GNS client data file required for installation, [5-12](#)
 name resolution for, [5-12](#)
GNS client data file
 how to create, [5-13](#)
GNS virtual IP address, [1-5](#)
grid home
 unlocking, [10-10](#)
grid infrastructure management repository, [9-3](#)
Grid Infrastructure Management Repository
 global, [8-10](#)
 local, [8-10](#)
Grid user
 creating, [6-15](#)

gridSetup script, [9-7](#), [9-14](#), [9-21](#)
groups
 creating the asmdba group, [6-12](#)
 creating the asmoper group, [6-13](#)
 creating the backupdba group, [6-14](#)
 creating the dba group, [6-13](#)
 creating the dgdba group, [6-14](#)
 creating the kmdba group, [6-14](#)
 creating the Oracle ASM group, [6-12](#)
 creating the racdba group, [6-14](#)
OINSTALL, [6-3](#)
OINSTALL group, [1-4](#)
OSBACKUPDBA (backupdba), [6-9](#)
OSDBA (dba), [6-9](#)
OSDBA group (dba), [6-9](#)
OSDGDBA (dgdba), [6-9](#)
OSKMDBA (kmdba), [6-9](#)
OSOPER (oper), [6-9](#)
OSOPER group (oper), [6-9](#)

H

hardware requirements
 display, [1-1](#)
 IPMI, [1-1](#)
 local storage for Oracle homes, [1-1](#)
 network, [1-1](#)
 RAM, [1-1](#)
 tmp, [1-1](#)
highly available IP addresses (HAIP), [5-6](#), [5-7](#)
host names
 legal host names, [1-5](#)
Hub Nodes, [5-18](#), [5-19](#)
hugepages, [1-4](#)
Hugepages,, [4-6](#)

I

IDE disks
 device names, [B-8](#)
image
 install, [9-2](#)
image-based installation of Oracle Grid Infrastructure,
 [9-7](#), [9-14](#), [9-21](#)
inaccessible nodes
 upgrading, [11-16](#)
incomplete installations, [11-31](#)
init.ora
 and SGA permissions, [10-7](#)
installation
 cloning a Grid infrastructure installation to other
 nodes, [9-29](#)
 response files
 preparing, [A-3](#), [A-5](#)
 templates, [A-3](#)
 silent mode, [A-6](#)
installation planning, [1-1](#)

- installation types
 - and Oracle ASM, [8-3](#)
- installer screen
 - Specify NFS Locations for ASM Disk Groups, [7-8](#)
 - Storage Option Information, [7-8](#)
- installer screens
 - ASM Storage Option, [8-14](#)
 - Cluster Node Information, [5-21](#)
 - Grid Plug and Play Information, [5-11](#), [5-21](#)
 - Network Interface Usage, [5-19](#)
 - Node Selection screen, [11-13](#)
- installing Oracle Member Cluster, [9-21](#)
- interconnect, [1-5](#)
- interconnects
 - single interface, [5-7](#)
- interfaces
 - requirements for private interconnect, [5-3](#)
- IPMI
 - addresses not configurable by GNS, [6-29](#)
 - preinstallation tasks, [6-28](#)
- IPv4 requirements, [5-2](#), [5-8](#)
- IPv6 requirements, [5-2](#), [5-8](#)

J

- JDK requirements, [4-10](#)

K

- kernel parameters
 - changing, [B-19](#)
 - displaying, [B-19](#)
 - SUSE Linux, [B-20](#)
 - tcp and udp, [B-21](#)
- kernel parameters configuration, [B-17](#)
- kmdba group
 - creating, [6-14](#)
- Korn shell
 - default user startup file, [6-23](#)

L

- Leaf Nodes, [5-18](#), [5-19](#)
- legal host names, [1-5](#)
- licensing, [1-12](#)
- Linux kernel parameters, [B-17](#)
- login file, [6-23](#)
- lsdev command, [B-8](#)
- LVM
 - recommendations for Oracle ASM, [8-3](#)

M

- Management Database, [8-10](#)
- management repository service, [9-3](#)
- manifest file, [8-21](#)
- mask

- mask (*continued*)
 - setting default file mode creation mask, [6-22](#)
- mixed binaries, [4-10](#)
- mode
 - setting default file mode creation mask, [6-22](#)
- multipath disks, [B-12](#)
- Multiple Oracle Homes Support
 - advantages, [C-2](#)
- multiversioning, [C-2](#)

N

- Name Service Cache Daemon
 - enabling, [4-30](#)
- Net Configuration Assistant (NetCA)
 - response files, [A-9](#)
 - running at command prompt, [A-9](#)
- netca.rsp file, [A-4](#)
- network interface cards
 - requirements, [5-6](#)
- network requirements, [1-5](#)
- network, minimum requirements, [1-1](#)
- networks
 - configuring interfaces, [5-25](#)
 - for Oracle Flex Clusters, [5-18](#), [5-19](#)
 - hardware minimum requirements, [5-6](#)
 - IP protocol requirements for, [5-2](#), [5-8](#)
 - manual address configuration example, [5-24](#)
 - Oracle Flex ASM, [1-5](#)
 - required protocols, [5-6](#)
- NFS
 - and data files, [7-10](#)
 - buffer size requirements, [8-17](#)
 - for data files, [7-10](#)
- NFS mounts
 - Direct NFS Client
 - requirements, [7-11](#)
 - mtab, [7-11](#)
 - oranfstab, [7-11](#)
- noninteractive mode
 - See response file mode

O

- OCR
 - See Oracle Cluster Registry
- OFA, [C-1](#)
 - See also Optimal Flexible Architecture
- oifcfg, [5-3](#)
- OINSTALL directory, [C-5](#)
- OINSTALL group
 - creating the oraInventory group, [6-3](#)
 - system privileges granted by, [6-3](#)
- OINSTALL group1, [1-7](#)
 - See also Oracle Inventory directory
- Opatch, [11-4](#)
- OPatch, [10-10](#)

- oper group
 - description, [6-9](#)
- operating system
 - different on cluster members, [4-10](#)
 - parameters, [B-17](#)
 - requirements, [4-10](#)
- operating system privileges groups, [1-7](#)
- operating system requirements, [1-2](#)
- Optimal Flexible Architecture
 - about, [C-1](#)
- ORAchk
 - and Upgrade Readiness Assessment, [1-12](#)
- Oracle ACFS
 - Installing Oracle RAC binaries not supported on Oracle Flex Cluster, [7-5](#)
 - restrictions for Oracle Restart, [7-5](#)
 - supported Linux x86-64 kernels, [7-4](#)
- Oracle ASM
 - characteristics of failure groups, [8-3, 8-15](#)
 - disk groups, [8-3](#)
 - failure groups
 - examples, [8-3](#)
 - identifying, [8-3](#)
 - guidelines, [7-7](#)
 - installing, [9-7, 9-14, 9-21](#)
 - recommendations for disk groups, [8-3](#)
 - space required for preconfigured database, [8-3](#)
- Oracle ASM Filter Driver
 - about, [8-14](#)
- Oracle ASM Filter Driver (Oracle ASMFDD), [9-14, 9-21](#)
- Oracle ASM group
 - creating, [6-12](#)
- Oracle ASM library driver (oracleasm)
 - installing, [B-6](#)
- Oracle ASM on NFS, [7-8](#)
- Oracle ASM password file, [7-7](#)
- Oracle ASMLIB
 - about, [B-5](#)
 - configuring, [B-5](#)
 - deinstalling on Oracle Grid Infrastructure, [B-14](#)
- Oracle Automatic Storage Management
 - checking disk availability, [B-8](#)
 - displaying attached disks, [B-8](#)
 - identifying available disks, [B-8](#)
 - identifying disks, [B-8](#)
 - part of Oracle Grid Infrastructure installation, [8-11](#)
 - performance, [4-30](#)
- Oracle Automatic Storage Management library driver
 - See Oracle ASMLIB
- Oracle base, [C-1, C-5](#)
- Oracle base directory, [6-7](#)
- Oracle Cluster Registry
 - configuration of, [1-9](#)
- Oracle Clusterware Files
 - NFS to Oracle ASM, [11-9](#)
- Oracle Database (*continued*)
 - data file storage options, [7-6](#)
 - requirements with Oracle ASM, [8-3](#)
- Oracle Database Configuration Assistant
 - response file, [A-4](#)
- Oracle Disk Manager
 - and Direct NFS, [8-20](#)
- Oracle Domain Services Cluster, [9-3, 9-5](#)
- Oracle Enterprise Manager, [11-20](#)
- Oracle Extended Cluster
 - converting to, [11-32](#)
- Oracle Extended Clusters, [9-6](#)
- Oracle Flex ASM
 - and Oracle ASM clients, [1-5](#)
 - networks, [1-5](#)
- Oracle Flex Clusters
 - about, [5-18](#)
 - and Hub Nodes, [10-9](#)
 - and Leaf Nodes, [10-9](#)
 - and Oracle Flex ASM, [5-19](#)
 - and Oracle Flex ASM cluster, [5-18](#)
- Oracle Restart
 - restrictions for, [7-5](#)
 - restrictions for Oracle ACFS, [7-5](#)
- Oracle Grid Infrastructure
 - restrictions for Oracle ACFS, [7-5](#)
- Oracle Grid Infrastructure upgrading, [11-17](#)
- Oracle home
 - ASCII path restriction for, [1-4](#)
 - file path, [C-5](#)
 - Grid home
 - filepath, [C-5](#)
 - naming conventions, [C-4](#)
- Oracle home directory, [6-7](#)
- Oracle Inventory
 - identifying existing, [6-2](#)
- Oracle Inventory Directory
 - OINSTALL group, [C-3](#)
- Oracle Inventory group
 - creating, [6-3](#)
- Oracle IO Server, [5-19](#)
- Oracle Layered File System, [9-31, 11-17](#)
- Oracle Member Clusters
 - for applications, [9-5](#)
 - for databases, [9-5](#)
- Oracle Net Configuration Assistant
 - response file, [A-4](#)
- Oracle Optimal Flexible Architecture
 - See Optimal Flexible Architecture
- Oracle Software Owner user
 - creating, [6-5, 6-15](#)
- Oracle Software Owner users
 - configuring environment for, [6-22](#)
 - determining default shell, [6-23](#)
- Oracle Standalone Cluster
 - storage option, [7-8](#)
- Oracle Standalone Clusters, [9-3](#)

- Oracle Universal Installer
 - response files
 - list of, [A-4](#)
- Oracle Upgrade Companion, [4-7](#)
- oracle user
 - creating, [6-5](#)
- Oracle user
 - configuring environment for, [6-22](#)
 - determining default shell, [6-23](#)
 - modifying, [6-15](#)
- ORACLE_BASE environment variable
 - removing from shell startup file, [6-23](#)
- ORACLE_HOME environment variable
 - removing from shell startup file, [6-23](#)
- ORACLE_SID environment variable
 - removing from shell startup file, [6-23](#)
- oracle-rdbms-server preinstallation RPM, [3-2](#)
- oracleasm RPM
 - installing, [B-6](#)
- oraInventory, [6-3](#), [C-5](#)
 - See also* Oracle Inventory group
- oranfstab configuration file, [8-18](#)
- oranfstab file, [8-20](#)
- OSASM group
 - creating, [6-12](#)
- OSBACKUPDBA group
 - creating, [6-14](#)
- OSBACKUPDBA group (backupdba), [6-9](#)
- OSDBA, [1-7](#)
- OSDBA for ASM
 - creating for Oracle Grid Infrastructure, [6-12](#)
- OSDBA groups
 - creating, [6-13](#)
 - creating for Oracle Grid Infrastructure, [6-13](#)
 - description for database, [6-9](#)
 - SYSDBA privilege, [6-9](#)
- OSDGDBA group
 - creating, [6-14](#)
- OSDGDBA group (dgdba), [6-9](#)
- OSKMDBA group
 - creating, [6-14](#)
- OSKMDBA group (kmdba), [6-9](#)
- OSOPER group
 - creating, [6-13](#)
- OSOPER groups
 - description for database, [6-9](#)
 - SYSOPER privilege, [6-9](#)
- OSRACDBA group
 - creating, [6-14](#)

P

- packages
 - oracle-rdbms-server preinstallation package, [3-2](#)
- parameter file
 - and permissions to read and write the SGA, [10-7](#)
- partition

- partition (*continued*)
 - using with Oracle ASM, [8-3](#)
- patch updates, [10-2](#)
- policy-managed databases
 - and SCAN, [5-5](#)
- postinstallation
 - recommended tasks
 - root.sh script, backing up, [10-3](#)
- postinstallation -executeConfigTools option, [A-10](#)
- postinstallation configToolAllCommands script, [A-13](#)
- preconfigured database
 - Oracle ASM disk space requirements, [8-3](#)
 - requirements when using Oracle ASM, [8-3](#)
- primary host name, [1-5](#)
- profile file, [6-23](#)
- proxy realm, [1-12](#)
- public node name
 - and primary host name, [1-5](#)

Q

- quorum disks
 - NFS mount, [7-8](#)

R

- racdba group
 - creating, [6-14](#)
- RAID
 - device names, [B-8](#)
 - recommended Oracle ASM redundancy level, [8-3](#)
- Rapid Home Provisioning, [9-31](#), [11-17](#)
- Rapid Home Provisioning Server, [9-3](#)
- recommendations
 - client access to the cluster, [10-5](#)
 - number of IP address for SCAN resolution, [5-5](#)
 - private network, [5-3](#)
- redundancy level
 - and space requirements for preconfigured database, [8-3](#)
- Redundant Interconnect Usage
 - IPv4 requirement, [5-7](#)
- registering resources, [11-20](#)
- releases
 - multiple, [C-2](#)
- relinking Oracle Grid Infrastructure home binaries, [10-10](#), [12-12](#)
- removing Oracle software
 - examples, [12-6](#)
- requiremenets
 - interconnects, [5-7](#)
- requirements
 - for networks, [5-6](#)
- resource parameters, [B-17](#)
- response file, [A-8](#)
- response file installation
 - preparing, [A-3](#)

- response file installation (*continued*)
 - response files
 - templates, [A-3](#)
 - silent mode, [A-6](#)
- response file mode
 - about, [A-2](#)
 - reasons for using, [A-2](#)
 - See also* response files, silent mode
- response files
 - about, [A-2](#)
 - creating with template, [A-4](#)
 - database configuration assistant, [A-8](#)
 - dbca.rsp, [A-4](#)
 - enterprise.rsp, [A-4](#)
 - general procedure, [A-3](#)
 - Net Configuration Assistant, [A-9](#)
 - netca.rsp, [A-4](#)
 - passing values at command line, [A-2](#)
 - specifying with Oracle Universal Installer, [A-6](#)
 - See also* silent mode.
- root user
 - logging in as, [2-1](#)
- root.sh script
 - backing up, [10-3](#)
- rootcrs.pl
 - restriction for Oracle Flex Cluster deinstallation, [12-13](#)
- rootcrs.sh, [12-2](#)
- roothas.sh, [12-2](#)
- RPMS
 - for completing operating system configuration, [4-4](#)
- running gridSetup.sh, [11-13](#)
- running multiple Oracle releases, [C-2](#)

S

- SCAN
 - required for clients of policy-managed databases, [5-5](#)
 - understanding, [5-5](#)
- SCAN address, [1-5](#)
- SCAN addresses, [5-5](#)
- SCAN listeners, [5-5](#)
- SCANs
 - client access, [10-5](#)
 - configuring, [1-5](#)
 - description, [10-5](#)
- SCSI disks
 - device names, [B-8](#)
- SELinux
 - restrictions on Oracle ACFS, [7-4](#)
- shell
 - determining default shell for Oracle user, [6-23](#)
- SHELL environment variable
 - checking value of, [6-23](#)
- shell startup file

- shell startup file (*continued*)
 - editing, [6-23](#)
 - removing environment variables, [6-23](#)
- silent mode
 - about, [A-2](#)
 - reasons for using, [A-2](#)
- silent mode installation, [A-6](#)
- single client access names
 - See* SCAN addresses
- software requirements, [4-10](#)
- space requirements, [8-7](#)
- ssh
 - and X11 Forwarding, [6-26](#)
 - configuring, [B-1](#)
- Standard cluster
 - upgrades result in, [11-4](#)
- standard operating environment, [9-31](#)
- startup file
 - for shell, [6-23](#)
- stty
 - suppressing to prevent installation errors, [6-27](#)
- swap space
 - allocation, [1-4](#)
- switches
 - minimum speed, [5-6](#)
- SYSBACKUPDBA system privileges, [6-9](#)
- SYSDBA privilege
 - associated group, [6-9](#)
- SYSDGDBA system privileges, [6-9](#)
- SYSKMDBA system privileges, [6-9](#)
- SYSOPER privilege
 - associated group, [6-9](#)
- system global area
 - permissions to read and write, [10-7](#)
- system privileges
 - SYSBACKUPDBA, [6-9](#)
 - SYSDGDBA, [6-9](#)
 - SYSKMDBA, [6-9](#)
- system requirements, [1-1](#)

T

- tcp_rmem, [8-17](#)
- tcp_wmem, [8-17](#)
- TCP/IP, [5-6](#)
- TEMP environment variable
 - commands
 - env, [6-23](#)
 - env command, [6-23](#)
 - environment
 - checking settings, [6-23](#)
 - setting, [6-23](#)
 - umask, [6-23](#)
 - umask command, [6-23](#)
- terminal output commands
 - suppressing for Oracle installation owner accounts, [6-27](#)

- TMPDIR environment variable
 - setting, [6-23](#)
- token-rings
 - unsupported, [5-6](#)
- Transparent Hugepages
 - disable for Oracle Database servers, [4-6](#), [4-29](#)
- troubleshooting
 - cron jobs and installation, [1-12](#)
 - disk space errors, [1-4](#)
 - environment path errors, [1-4](#)
 - garbage strings in script inputs found in log files, [6-27](#)
 - I/O lags, [4-29](#)
 - inventory corruption, [6-15](#)
 - nfs mounts, [4-30](#)
 - ocssd.log threads blocked, [4-29](#)
 - public network failures, [4-30](#)
 - root.sh errors, [12-13](#)
 - ssh, [B-2](#), [B-4](#)
 - ssh errors, [6-27](#)
 - stty errors, [6-27](#)
 - unconfiguring Oracle Clusterware to fix causes of root.sh errors, [12-13](#)
 - unset environment variables, [1-4](#)
 - user equivalency, [B-2](#), [B-4](#)
 - user equivalency errors, [6-3](#)
- Troubleshooting
 - DBCA does not recognize Oracle ASM disk size and fails to create disk groups, [10-10](#)
- typographic conventions, [xviii](#)

U

- umask command, [6-22](#)
- unconfiguring Oracle Clusterware, [12-13](#)
- uninstall
 - See removing Oracle software
- UNIX commands
 - fdisk, [B-8](#)
 - lsdev, [B-8](#)

- UNIX commands (*continued*)
 - xhost, [2-1](#)
- UNIX workstation
 - installing from, [2-1](#)
- unreachable nodes
 - upgrading, [11-16](#)
- upgrade
 - Oracle Automatic Storage Management, [8-11](#)
 - running gridSetup.sh, [11-13](#)
- upgrade tasks, [11-20](#)
- upgrades
 - and Leaf Nodes, [11-4](#)
 - and SCAN, [5-5](#)
 - best practices, [4-7](#)
 - restrictions for, [11-4](#)
 - unsettling environment variables for, [11-8](#)
- upgrading
 - and ORAchk Upgrade Readiness Assessment, [1-12](#)
 - inaccessible nodes, [11-16](#)
 - options, [4-8](#)
- useradd command, [6-15](#)
- users
 - creating the oracle user, [6-5](#)

V

- vendor clusterware
 - and cluster names for Oracle Grid Infrastructure, [1-5](#)
- voting files
 - configuration of, [1-9](#)

X

- X Window System
 - enabling remote hosts, [2-1](#)
- X11 forwarding errors, [6-26](#), [B-4](#)
- xhost command, [2-1](#)
- xtitle
 - suppressing to prevent installation errors, [6-27](#)