

Oracle® Grid Infrastructure

Installation Guide

12c Release 1 (12.1) for Linux

E48914-22

July 2017

Oracle Grid Infrastructure Installation Guide, 12c Release 1 (12.1) for Linux

E48914-22

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

Primary Author: Aparna Kamath

Contributing Authors: Douglas Williams, Gavin Bowe, David Brower, Jonathan Creighton, Paul Harter, Prakash Jashnani, Christopher Jones, Markus Michalewicz, Janet Stern, Rich Strohm, Rick Wessman

Contributor: The Database 12c documentation is dedicated to Mark Townsend, who was an inspiration to all who worked on this release.

Contributors: Harshit Agrawal, Prasad Bagal, Subhransu Basu, Mark Bauer, Saar Maoz, Eric Belden, Barb Glover, Donald Graves, Eugene Karichkin, Aneesh Khandelwal, Reema Khosla, Virkumar Koganole, Erich Kreisler, Jai Krishnani, Sergio Leunissen, John Leys, Richard Long, Rudregowda Mallegowda, Saar Maoz, Mughees Minhas, John McHugh, Balaji Pagadala, Srinivas Poovala, Gurumurthy Ramamurthy, Sunil Ravindrachar, Vipin Samar, Mark Scardina, Santhosh Selvaraj, Cathy Shea, Malaiarasan Stalin, Binoy Sukumaran, Roy Swonger, Randy Urbano, Preethi Vallam, Rui Wang, Jim Williams, Yíyan Yang

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	xiii
Intended Audience.....	xiii
Documentation Accessibility	xiii
Related Documents	xiii
Conventions	xv
Changes in This Release for Oracle Grid Infrastructure Installation Guide	xvii
Changes in Oracle Grid Infrastructure 12c Release 1 (12.1)	xvii
1 Oracle Grid Infrastructure Installation Checklist	
1.1 System Hardware, Software and Configuration Checklists.....	1-1
1.1.1 Oracle Grid Infrastructure Installation Server Hardware Checklist.....	1-1
1.1.2 Oracle Grid Infrastructure and Oracle RAC Environment Checklist	1-2
1.1.3 Oracle Grid Infrastructure Network Checklist.....	1-3
1.1.4 Oracle Grid Infrastructure and Oracle RAC Upgrades Checklist	1-4
1.1.5 Oracle Grid Infrastructure Storage Configuration Tasks	1-6
1.1.6 Oracle Grid Infrastructure Starting the Installation Tasks	1-6
2 Configuring Servers for Oracle Grid Infrastructure and Oracle RAC	
2.1 Checking Server Hardware and Memory Configuration	2-1
2.2 General Server Minimum Requirements.....	2-2
2.3 Server Storage Minimum Requirements	2-2
2.4 Server Memory Minimum Requirements	2-3
2.4.1 Minimum Memory Requirements for Oracle Grid Infrastructure	2-3
2.4.2 Shared Memory Requirements	2-4
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 From Unbreakable Linux Network.....	3-2
3.3 Installing Oracle Linux with Oracle Linux Yum Server Support	3-3
3.4 Installing the Oracle Preinstallation RPM From DVDs or Images	3-4
3.5 Additional Optional Operating System Configuration Tasks.....	3-4
3.5.1 Configure Oracle Ksplice Repository for Oracle Linux	3-5
3.5.2 Configure Additional Operating System Features	3-5

3.6	Required System Configuration for Oracle Grid Infrastructure.....	3-5
-----	---	-----

4 Configuring Operating Systems for Oracle Grid Infrastructure and Oracle RAC

4.1	Guidelines for Linux Operating System Installation.....	4-1
4.1.1	Completing a Minimal Linux Installation.....	4-2
4.1.2	Completing a Default Linux Installation.....	4-3
4.1.3	About Unbreakable Enterprise Kernel for Oracle Linux	4-3
4.1.4	About the Oracle Preinstallation RPM	4-4
4.1.5	Using Oracle Ksplice to Perform a Zero Downtime Update.....	4-5
4.2	Reviewing Operating System and Software Upgrade Best Practices	4-6
4.2.1	General Upgrade Best Practices.....	4-6
4.2.2	Oracle ASM Upgrade Notifications	4-6
4.2.3	Rolling Upgrade Procedure Notifications.....	4-7
4.3	Reviewing Operating System Security Common Practices	4-7
4.4	Using Installation Fixup Scripts.....	4-7
4.5	Logging In to a Remote System Using X Terminal.....	4-8
4.6	Using Oracle RPM Checker on IBM: Linux on System z	4-9
4.7	About Operating System Requirements.....	4-9
4.8	Operating System Requirements for x86-64 Linux Platforms	4-10
4.8.1	Supported Oracle Linux 7 and Red Hat Linux 7 Distributions for x86-64.....	4-11
4.8.2	Supported Oracle Linux 6 and Red Hat Linux 6 Distributions for x86-64.....	4-12
4.8.3	Supported Oracle Linux 5 and Red Hat Linux 5 Distributions for x86-64.....	4-13
4.8.4	Supported SUSE Linux Enterprise Server Distributions for x86-64.....	4-14
4.8.5	Supported NeoKylin Linux Advanced Server Distributions for x86-64	4-16
4.9	Operating System Requirements for IBM: Linux on System z.....	4-17
4.9.1	Supported Red Hat Enterprise Linux 7 Distributions for IBM: Linux on System z	4-18
4.9.2	Supported Red Hat Enterprise Linux 6 Distributions for IBM: Linux on System z	4-18
4.9.3	Supported Red Hat Enterprise Linux 5 Distributions for IBM: Linux on System z	4-19
4.9.4	Supported SUSE Distributions for IBM: Linux on System z	4-20
4.10	Additional Drivers and Software Packages for Linux.....	4-22
4.10.1	Installation Requirements for Open Database Connectivity.....	4-23
4.10.2	Installation Requirements for PAM on Linux	4-23
4.10.3	Installation Requirements for OCFS2	4-23
4.10.4	Installation Requirements for Oracle Messaging Gateway	4-24
4.10.5	Installation Requirements for Lightweight Directory Access Protocol	4-24
4.10.6	Installation Requirements for Programming Environments for Linux	4-25
4.10.7	Installation Requirements for Web Browsers	4-26
4.11	Checking the Software Requirements.....	4-26
4.12	Installing the cvuqdisk RPM for Linux.....	4-27
4.13	Checking Shared Memory File System Mount on Linux.....	4-28
4.14	Enabling the Name Service Cache Daemon.....	4-28
4.15	Setting the Disk I/O Scheduler on Linux	4-29
4.16	Setting Network Time Protocol for Cluster Time Synchronization	4-29
4.17	Using Automatic SSH Configuration During Installation.....	4-30

5 Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

5.1	Network Interface Hardware Requirements	5-1
5.2	IP Interface Configuration Requirements	5-2
5.3	Private Interconnect Redundant Network Requirements	5-3
5.4	IPv4 and IPv6 Protocol Requirements	5-3
5.5	Oracle Grid Infrastructure IP Name and Address Requirements	5-4
5.5.1	About Oracle Grid Infrastructure Name Resolution Options.....	5-5
5.5.2	Cluster Name and SCAN Requirements.....	5-5
5.5.3	IP Name and Address Requirements For Grid Naming Service (GNS).....	5-6
5.5.4	IP Name and Address Requirements For Multi-Cluster GNS.....	5-6
5.5.5	IP Name and Address Requirements for Standard Cluster Manual Configuration .	5-8
5.5.6	Confirming the DNS Configuration for SCAN.....	5-9
5.6	About Oracle Flex ASM Clusters Networks	5-9
5.7	Broadcast Requirements for Networks Used by Oracle Grid Infrastructure.....	5-11
5.8	Multicast Requirements for Networks Used by Oracle Grid Infrastructure.....	5-11
5.9	Domain Delegation to Grid Naming Service.....	5-11
5.9.1	Choosing a Subdomain Name for Use with Grid Naming Service.....	5-11
5.9.2	Configuring DNS for Cluster Domain Delegation to Grid Naming Service	5-12
5.10	Configuration Requirements for Oracle Flex Clusters	5-13
5.10.1	General Requirements for Oracle Flex Cluster Configuration.....	5-13
5.10.2	Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses.....	5-13
5.10.3	Oracle Flex Cluster Manually-Assigned Addresses	5-13
5.11	Grid Naming Service Standard Cluster Configuration Example	5-14
5.12	Manual IP Address Configuration Example.....	5-15
5.13	Network Interface Configuration Options.....	5-16
5.14	Multiple Private Interconnects and Oracle Linux	5-17

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

6.1	Creating Groups, Users and Paths for Oracle Grid Infrastructure.....	6-1
6.1.1	Determining If the Oracle Inventory and Oracle Inventory Group Exists.....	6-1
6.1.2	Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist	6-2
6.1.3	Creating the Oracle Grid Infrastructure User.....	6-3
6.1.4	About the Oracle Base Directory for the Grid User.....	6-6
6.1.5	About the Oracle Home Directory for Oracle Grid Infrastructure Software	6-6
6.1.6	Creating the Oracle Home and Oracle Base Directory.....	6-7
6.1.7	About Job Role Separation Operating System Privileges Groups and Users	6-8
6.1.8	Descriptions of Job Role Separation Groups and Users.....	6-9
6.1.9	Creating Job Role Separation Operating System Privileges Groups and User.....	6-12
6.1.10	Example of Creating Minimal Groups, Users, and Paths	6-16
6.1.11	Example of Creating Role-allocated Groups, Users, and Paths.....	6-18
6.2	Configuring Grid Infrastructure Software Owner User Environments	6-20
6.2.1	Environment Requirements for Oracle Software Owners.....	6-21
6.2.2	Procedure for Configuring Oracle Software Owner Environments	6-21
6.2.3	Checking Resource Limits for the Oracle Software Installation Users	6-23
6.2.4	Setting Remote Display and X11 Forwarding Configuration	6-24

6.2.5	Preventing Installation Errors Caused by Terminal Output Commands	6-25
6.3	Enabling Intelligent Platform Management Interface (IPMI).....	6-25
6.3.1	Requirements for Enabling IPMI.....	6-26
6.3.2	Configuring the IPMI Management Network.....	6-26
6.3.3	Configuring the IPMI Driver	6-26
6.4	Determining Root Script Execution Plan.....	6-31

7 Configuring Storage for Oracle Grid Infrastructure and Oracle RAC

7.1	Reviewing Oracle Grid Infrastructure Storage Options.....	7-1
7.1.1	Supported Storage Options	7-1
7.1.2	About Oracle ACFS and Oracle ADVM.....	7-3
7.1.3	General Storage Considerations for Oracle Grid Infrastructure and Oracle RAC....	7-5
7.1.4	Guidelines for Using Oracle ASM Disk Groups for Storage	7-6
7.1.5	Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC	7-6
7.1.6	After You Have Selected Disk Storage Options	7-7
7.2	About Shared File System Storage Configuration	7-7
7.2.1	Guidelines for Using a Shared File System with Oracle Grid Infrastructure	7-7
7.2.2	Requirements for Oracle Grid Infrastructure Shared File System Volume Sizes.....	7-8
7.2.3	Deciding to Use a Cluster File System for Oracle Clusterware Files	7-10
7.2.4	About Direct NFS Client and Data File Storage	7-10
7.2.5	Deciding to Use NFS for Data Files.....	7-12
7.3	Configuring Operating System and Direct NFS Client.....	7-12
7.3.1	Configuring Operating System NFS Mount and Buffer Size Parameters	7-13
7.3.2	Checking Operating System NFS Mount and Buffer Size Parameters	7-13
7.3.3	Checking NFS Mount and Buffer Size Parameters for Oracle RAC.....	7-14
7.3.4	Checking TCP Network Protocol Buffer for Direct NFS Client.....	7-15
7.3.5	Enabling Direct NFS Client Oracle Disk Manager Control of NFS.....	7-15
7.3.6	Enabling Hybrid Columnar Compression on Direct NFS Client	7-17
7.3.7	Specifying Network Paths with the Oranfstab File	7-17
7.3.8	Creating Directories for Oracle Clusterware Files on Shared File Systems	7-18
7.3.9	Creating Directories for Oracle Database Files on Shared File Systems.....	7-19
7.3.10	Disabling Direct NFS Client Oracle Disk Management Control of NFS	7-20
7.4	Oracle Automatic Storage Management Storage Configuration	7-20
7.4.1	Configuring Storage for Oracle Automatic Storage Management	7-20
7.4.2	About Oracle ASM with Oracle ASM Filter Driver.....	7-27
7.4.3	Using Disk Groups with Oracle Database Files on Oracle ASM.....	7-28
7.4.4	Configuring Oracle Automatic Storage Management Cluster File System	7-29
7.4.5	Upgrading Existing Oracle ASM Instances	7-30
7.5	Configuring Raw Logical Volumes on IBM: Linux on System z	7-31

8 Installing Oracle Grid Infrastructure for a Cluster

8.1	Installing Oracle Grid Infrastructure	8-1
8.1.1	Running OUI to Install Oracle Grid Infrastructure	8-1
8.1.2	Installing Oracle Grid Infrastructure Using a Cluster Configuration File	8-4
8.2	Installing Grid Infrastructure Using a Software-Only Installation	8-4
8.2.1	Installing the Software Binaries	8-5
8.2.2	Configuring the Software Binaries	8-5

8.2.3	Configuring the Software Binaries Using a Response File	8-6
8.2.4	Setting Ping Targets for Network Checks	8-6
8.3	Confirming Oracle Clusterware Function	8-7
8.4	Confirming Oracle ASM Function for Oracle Clusterware Files	8-7
8.5	Understanding Offline Processes in Oracle Grid Infrastructure	8-8

9 Oracle Grid Infrastructure Postinstallation Procedures

9.1	Required Postinstallation Tasks	9-1
9.2	Recommended Postinstallation Tasks	9-2
9.2.1	Tuning Semaphore Parameters	9-2
9.2.2	Create a Fast Recovery Area Disk Group	9-3
9.2.3	Checking the SCAN Configuration	9-4
9.2.4	Downloading and Installing the ORAchk Health Check Tool	9-4
9.2.5	Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications 9-5	
9.3	Using Earlier Oracle Database Releases with Oracle Grid Infrastructure	9-5
9.3.1	General Restrictions for Using Earlier Oracle Database Versions	9-5
9.3.2	Managing Server Pools with Earlier Database Versions	9-6
9.3.3	Making Oracle ASM Available to Earlier Oracle Database Releases	9-6
9.3.4	Using ASMCA to Administer Disk Groups for Earlier Database Versions	9-7
9.3.5	Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x	9-7
9.3.6	Using the Correct LSNRCTL Commands	9-8
9.4	Modifying Oracle Clusterware Binaries After Installation	9-8

10 How to Modify or Deinstall Oracle Grid Infrastructure

10.1	Deciding When to Deinstall Oracle Clusterware	10-1
10.2	Migrating Standalone Grid Infrastructure Servers to a Cluster	10-2
10.3	Relinking Oracle Grid Infrastructure for a Cluster Binaries	10-3
10.4	Changing the Oracle Grid Infrastructure Home Path	10-4
10.5	Unconfiguring Oracle Clusterware Without Removing Binaries	10-5
10.6	Removing Oracle Clusterware and Oracle ASM	10-6
10.6.1	About the Deinstallation Tool	10-6
10.6.2	Deinstallation Tool Command Example for Oracle Grid Infrastructure	10-9
10.6.3	Deinstallation Response File Example for Grid Infrastructure for a Cluster	10-10

A Troubleshooting the Oracle Grid Infrastructure Installation Process

A.1	Best Practices for Contacting Oracle Support	A-1
A.2	General Installation Issues	A-2
A.2.1	Other Installation Issues and Errors	A-7
A.3	Interpreting CVU "Unknown" Output Messages Using Verbose Mode	A-7
A.4	Interpreting CVU Messages About Oracle Grid Infrastructure Setup	A-8
A.5	About the Oracle Clusterware Alert Log	A-10
A.6	Missing Operating System Packages On Linux	A-10
A.7	Performing Cluster Diagnostics During Oracle Grid Infrastructure Installations	A-11
A.8	About Using CVU Cluster Healthchecks After Installation	A-11
A.9	Interconnect Configuration Issues	A-12

A.10	SCAN VIP and SCAN Listener Issues	A-13
A.11	Storage Configuration Issues	A-13
A.11.1	Recovery from Losing a Node Filesystem or Grid Home	A-14
A.11.2	Oracle ASM Library Driver Issues	A-14
A.11.3	Oracle ASM Issues After Upgrading Oracle Grid Infrastructure.....	A-15
A.11.4	Oracle ASM Issues After Downgrading Oracle Grid Infrastructure for Standalone Server (Oracle Restart) A-15	
A.12	Failed or Incomplete Installations and Upgrades	A-16
A.12.1	Completing Failed or Interrupted Upgrades.....	A-17
A.12.2	Completing Failed or Interrupted Installations	A-18

B How to Upgrade to Oracle Grid Infrastructure 12c Release 1

B.1	Back Up the Oracle Software Before Upgrades.....	B-1
B.2	About Oracle Grid Infrastructure and Oracle ASM Upgrade and Downgrade.....	B-2
B.3	Options for Oracle Grid Infrastructure Upgrades and Downgrades.....	B-2
B.4	Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades.....	B-3
B.5	Preparing to Upgrade an Existing Oracle Clusterware Installation.....	B-5
B.5.1	Checks to Complete Before Upgrading Oracle Clusterware.....	B-5
B.5.2	Unset Oracle Environment Variables	B-6
B.5.3	Running the Oracle ORAchK Upgrade Readiness Assessment	B-6
B.6	Using CVU to Validate Readiness for Oracle Clusterware Upgrades	B-7
B.6.1	About the CVU Grid Upgrade Validation Command Options	B-7
B.6.2	Example of Verifying System Upgrade Readiness for Grid Infrastructure	B-8
B.7	Understanding Rolling Upgrades Using Batches	B-8
B.8	Performing Rolling Upgrade of Oracle Grid Infrastructure.....	B-9
B.8.1	Performing a Standard Upgrade from an Earlier Release	B-9
B.8.2	Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable	B-10
B.8.3	Upgrading Inaccessible Nodes After Forcing an Upgrade.....	B-11
B.8.4	Changing the First Node for Install and Upgrade.....	B-11
B.9	Restrictions and Guidelines for Upgrading and Patching Oracle ASM.....	B-11
B.10	Performing Rolling Upgrade of Oracle ASM.....	B-12
B.10.1	Upgrading Oracle ASM Using ASMCA.....	B-12
B.11	Applying Patches to Oracle ASM.....	B-13
B.11.1	About Individual (One-Off) Oracle ASM Patches	B-13
B.11.2	About Oracle ASM Software Patch Levels	B-13
B.11.3	Patching Oracle ASM to a Software Patch Level	B-14
B.12	Updating Oracle Enterprise Manager Cloud Control Target Parameters.....	B-14
B.12.1	Updating the Enterprise Manager Cloud Control Target After Upgrades.....	B-14
B.12.2	Updating the Enterprise Manager Agent Base Directory After Upgrades	B-15
B.13	Unlocking the Existing Oracle Clusterware Installation.....	B-15
B.14	Checking Cluster Health Monitor Repository Size After Upgrading.....	B-16
B.15	Downgrading Oracle Clusterware After an Upgrade	B-16
B.15.1	About Downgrading Oracle Clusterware After an Upgrade.....	B-16
B.15.2	Downgrading to Releases Before 11g Release 2 (11.2.0.2).....	B-17
B.15.3	Downgrading to 11g Release 1 (11.2.0.2) or Later Release.....	B-18

C Installing and Configuring Oracle Database Using Response Files

C.1	How Response Files Work.....	C-1
C.1.1	Reasons for Using Silent Mode or Response File Mode	C-2
C.1.2	General Procedure for Using Response Files	C-2
C.2	Preparing a Response File.....	C-3
C.2.1	Editing a Response File Template	C-3
C.2.2	Recording a Response File.....	C-4
C.3	Running the Installer Using a Response File	C-5
C.4	Running Net Configuration Assistant Using a Response File	C-6
C.5	Postinstallation Configuration Using a Response File	C-6
C.5.1	About the Postinstallation Configuration File.....	C-7
C.5.2	Running Postinstallation Configuration Using a Response File.....	C-7

D Configuring Large Memory Optimization

D.1	Overview of HugePages	D-1
D.1.1	What HugePages Provides.....	D-1
D.2	Restrictions for HugePage Configurations	D-2
D.3	Disabling Transparent HugePages.....	D-2

E Oracle Grid Infrastructure for a Cluster Installation Concepts

E.1	Understanding Preinstallation Configuration.....	E-1
E.1.1	Optimal Flexible Architecture Guidelines for Oracle Grid Infrastructure.....	E-1
E.1.2	Oracle Grid Infrastructure for a Cluster and Oracle Restart Differences	E-2
E.1.3	Understanding the Oracle Inventory Group	E-2
E.1.4	Understanding the Oracle Inventory Directory	E-3
E.1.5	Understanding the Oracle Base directory	E-3
E.1.6	Understanding the Oracle Home for Oracle Grid Infrastructure Software	E-4
E.1.7	Location of Oracle Base and Oracle Grid Infrastructure Software Directories	E-5
E.2	Understanding Network Addresses	E-5
E.2.1	About the Public IP Address.....	E-5
E.2.2	About the Private IP Address	E-5
E.2.3	About the Virtual IP Address	E-6
E.2.4	About the Grid Naming Service (GNS) Virtual IP Address.....	E-6
E.2.5	About the SCAN for Oracle Grid Infrastructure Installations	E-6
E.3	Understanding Network Time Requirements	E-8
E.4	Understanding Oracle Flex Clusters and Oracle ASM Flex Clusters.....	E-8
E.5	Understanding Storage Configuration	E-9
E.5.1	Understanding Oracle Automatic Storage Management Cluster File System	E-9
E.5.2	About Migrating Existing Oracle ASM Instances.....	E-9
E.5.3	Standalone Oracle ASM Installations to Clustered Installation Conversions	E-10
E.6	Understanding Out-of-Place Upgrade.....	E-10

F How to Complete Preinstallation Tasks Manually

F.1	Configuring SSH Manually on All Cluster Nodes.....	F-1
F.1.1	Checking Existing SSH Configuration on the System.....	F-2

F.1.2	Configuring SSH on Cluster Nodes	F-2
F.1.3	Enabling SSH User Equivalency on Cluster Nodes.....	F-4
F.2	Configuring Kernel Parameters	F-5
F.2.1	Minimum Parameter Settings for Installation	F-5
F.2.2	Additional Parameter and Kernel Settings for SUSE Linux Enterprise Server	F-6
F.3	Setting UDP and TCP Kernel Parameters Manually	F-7
F.4	Configuring Storage Paths and Disk Devices	F-7
F.4.1	Configuring Storage Device Path Persistence Using Oracle ASMLIB.....	F-8
F.4.2	Configuring Disk Devices Manually for Oracle ASM.....	F-16
F.5	Checking OCFS2 Version Manually.....	F-18

Index

List of Tables

1-1	Server Hardware Checklist for Oracle Grid Infrastructure	1-1
1-2	Environment Configuration for Oracle Grid Infrastructure and Oracle RAC.....	1-2
1-3	Network Configuration Tasks for Oracle Grid Infrastructure and Oracle RAC	1-4
1-4	Installation Upgrades Checklist for Oracle Grid infrastructure	1-5
1-5	Oracle Grid Infrastructure Storage Configuration Checks.....	1-6
1-6	Oracle Grid Infrastructure Checks to Perform Before Starting the Installer.....	1-6
2-1	Swap Space Required for 64-bit Linux and Linux on System z	2-3
4-1	x86-64 Linux 7 Minimum Operating System Requirements	4-11
4-2	x86-64 Linux 6 Minimum Operating System Requirements	4-12
4-3	x86-64 Linux 5 Minimum Operating System Requirements	4-14
4-4	x86-64 Supported SUSE Linux Enterprise Server Operating System Requirements	4-15
4-5	x86-64 Supported NeoKylin Linux Minimum Operating System Requirements	4-16
4-6	IBM: Linux on System z Linux 7 Minimum Operating System Requirements	4-18
4-7	IBM: Linux on System z Linux 6 Minimum Operating System Requirements	4-19
4-8	IBM: Linux on System z Linux 5 Minimum Operating System Requirements	4-20
4-9	IBM: Linux on System z SUSE Minimum Operating System Requirements.....	4-21
4-10	Requirements for Programming Environments for x86-64 Linux.....	4-25
4-11	Requirements for Programming Environments for IBM: Linux on System z.....	4-26
4-12	Requirements for Programming Environments for Linux on SPARC.....	4-26
5-1	Grid Naming Service Example Network.....	5-15
5-2	Manual Network Configuration Example	5-16
6-1	Installation Owner Resource Limit Recommended Ranges	6-23
7-1	Supported Storage Options for Oracle Clusterware and Oracle RAC.....	7-2
7-2	Platforms That Support Oracle ACFS and Oracle ADVM.....	7-3
7-3	Oracle Clusterware Shared File System Volume Size Minimum Requirements.....	7-8
7-4	Oracle RAC Shared File System Volume Size Minimum Requirements.....	7-9
7-5	Oracle Clusterware Minimum Storage Space Required by Redundancy Type	7-22
7-6	Total Oracle Database Storage Space Required by Redundancy Type.....	7-23
C-1	Response Files for Oracle Database.....	C-3
C-2	Response files for Oracle Grid Infrastructure.....	C-3
F-1	Minimum Operating System Parameter Settings for Installation on Linux	F-6
F-2	ORACLEASM Configure Prompts and Responses	F-10
F-3	Types of Linux Storage Disk Paths.....	F-11
F-4	Disk Management Tasks Using ORACLEASM.....	F-12

Preface

Oracle Grid Infrastructure Installation Guide for Linux 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.

Intended Audience

Oracle Grid Infrastructure Installation Guide for Linux 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 Documents

For more information, see the following Oracle resources:

Oracle Clusterware and Oracle Real Application Clusters Documentation

This installation guide provides the steps required to complete an Oracle Clusterware and Oracle Automatic Storage Management installation, and to perform preinstallation steps for Oracle RAC.

If you intend to install Oracle RAC, then complete Oracle RAC preinstallation tasks as described in this installation guide. For Oracle RAC or Oracle Database installations, refer to the installation guides for these products after you complete the Oracle Grid Infrastructure installation.

Installation Guides

- *Oracle Database Installation Guide for Linux*
- *Oracle Real Application Clusters Installation Guide for Linux and UNIX*

Operating System-Specific Administrative Guides

- *Oracle Database Administrator's Reference, 12c Release 1 (12.1) for UNIX Systems*

Oracle Clusterware and Oracle Automatic Storage Management Administrative Guides

- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Automatic Storage Management Administrator's Guide*

Oracle Real Application Clusters Administrative Guides

- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Enterprise Manager Real Application Clusters Guide Online Help*

Generic Documentation

- *Oracle Database 2 Day DBA*
- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database New Features Guide*
- *Oracle Database Net Services Administrator's Guide*
- *Oracle Database Reference*

Printed documentation is available for sale in the Oracle Store at the following website:

<https://shop.oracle.com>

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network. You must register online before using Oracle Technology Network; registration is free and can be done at the following website:

<https://support.oracle.com>

If you already have a username and password for Oracle Technology Network, then you can go directly to the documentation section of the Oracle Technology Network website:

<http://www.oracle.com/technetwork/indexes/documentation/index.html>

Oracle error message documentation is available only in HTML. You can browse the error messages by range in the Documentation directory of the installation media. When you find a range, use your browser's search feature to locate a specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

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 Installation Guide

This preface contains:

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

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.

- [New Features for Oracle Grid Infrastructure 12c Release 1 \(12.1.0.2\)](#)
- [New Features for Oracle Grid Infrastructure 12c Release 1 \(12.1.0.1\)](#)
- [Deprecated Features](#)
- [Desupported Features](#)
- [Other Changes](#)

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

- 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. Oracle ASM uses the filter driver to validate write I/O requests to Oracle ASM disks.

The Oracle ASM filter driver rejects any I/O requests that are invalid. This action eliminates accidental overwrites of Oracle ASM disks that would cause corruption in the disks and files within the disk group. For example, the Oracle ASM filter driver filters out all non-Oracle I/Os which could cause accidental overwrites.

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

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

- 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 Servers (RHPS) clusters provide gold images to Rapid Home Provisioning Clients (RHPC).

See *Oracle Clusterware Administration and Deployment Guide* and *Oracle Real Application Clusters Installation Guide*.

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

- 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).

Note: The Grid Infrastructure Management Repository is not available on IBM: Linux on System z.

- 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)

- Cluster Health Monitor Enhancements for Oracle Flex Cluster

Cluster Health Monitor (CHM) has been enhanced to provide a highly available server monitor service that provides improved detection of operating system and cluster resource-related degradation and failures. In addition, CHM supports Oracle Flex Cluster configurations, including the ability for data collectors to collect from every node of the cluster and provide a single cluster representation of the data.

See Also: See *Oracle Clusterware Administration and Deployment Guide*

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

- Oracle Flex Cluster

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. In order 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 [Section 5.10, "Configuration Requirements for Oracle Flex Clusters"](#)

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about Oracle Flex Clusters, and *Oracle Grid Infrastructure Installation Guide* for more information about Oracle Flex Cluster deployment

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

- Oracle Cluster Registry Backup in ASM Disk Group Support

The Oracle Cluster Registry (OCR) backup mechanism enables storing the OCR backup in an Oracle ASM disk group. Storing the OCR backup in an Oracle ASM disk group simplifies OCR management by permitting access to the OCR backup from any node in the cluster should an OCR recovery become necessary.

- 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 connection requests to the appropriate database listener for the IP protocol of the client request.

See [Section 5.4, "IPv4 and IPv6 Protocol Requirements"](#)

- Grid Infrastructure Script Automation for Installation and Upgrade

This feature enables running any script requiring `root` privileges through the installer and other configuration assistants, so that you are no longer required to run root-based scripts manually during deployment.

Using script automation for installation and upgrade eliminates the need to run scripts manually on each node during the final steps of an Oracle Grid Infrastructure installation or upgrade.

See [Section 1.1.2, "Oracle Grid Infrastructure and Oracle RAC Environment Checklist"](#)

- Oracle Grid Infrastructure Rolling Migration for One-Off Patches

Oracle Grid Infrastructure one-off patch rolling migration and upgrade for Oracle ASM and Oracle Clusterware enables you to independently upgrade or patch clustered Oracle Grid Infrastructure nodes with one-off patches, without affecting database availability. This feature provides greater uptime and patching flexibility. This release also introduces a new Cluster state, "Rolling Patch." Operations allowed in a patch quiesce state are similar to the existing "Rolling Upgrade" cluster state.

See [Appendix B.10, "Performing Rolling Upgrade of Oracle ASM"](#)

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about ASM rolling migrations and patches

- Oracle Flex ASM Server and Oracle CloudFS

Oracle Flex ASM decouples the Oracle ASM instance from database servers and enables the Oracle ASM instance to run on a separate physical server from the database servers. Any number of Oracle ASM instances can be clustered to support numerous database clients. This is a component feature of Oracle CloudFS.

Oracle CloudFS is a storage cloud infrastructure with resource pooling, network accessibility, rapid elasticity and rapid provisioning that are key requirements for cloud computing environments.

This feature enables you to consolidate all 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 Cluster Synchronization Services (CSS) cluster. Depending on the performance requirements, you can make policy decisions on how various Oracle ASM clients access its files in a disk group.

Oracle Flex ASM supports Oracle Database 12c Release 1 (12.1) and later. Oracle Database 10g Release 2 (10.2) or later through Oracle Database 11g Release 2 (11.2) can continue to use ASM disk groups with no requirement to install patches.

See [Section 5.6, "About Oracle Flex ASM Clusters Networks"](#)

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about using Oracle Flex ASM servers

- 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* for more information about managing applications with policies

- Shared Grid Naming Service (GNS) Across Multiple Clusters

In previous releases, the Grid Naming Service (GNS) was dedicated to one Oracle Grid Infrastructure-based cluster, providing name resolution only for its own cluster member nodes. With this release, one Oracle GNS can now manage just the cluster member nodes in its own cluster, or GNS can provide naming resolution for all nodes across all clusters in the data center that are delegated to Oracle GNS for resolution.

Using only one Oracle GNS for all nodes that are part of an Oracle Grid Infrastructure cluster in the data center not only streamlines the naming convention, but also enables a data center cloud, minimizing day-to-day administration efforts.

See [Section 5.5, "Oracle Grid Infrastructure IP Name and Address Requirements"](#)

- Support for Separation of Database Administration Duties

Oracle Database 12c Release 1 (12.1) 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 [Section 6.1.8.3, "Extended Oracle Database Groups for Job Role Separation"](#)

See Also:

- *Oracle Database Administrator's Guide* for an overview of system privileges and operating system authentication
- *Oracle Database Security Guide* for information about using system privileges

Deprecated Features

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

- Change for Standalone Deinstallation tool
The deinstallation tool is now integrated with the installation media.
- Deprecation of `-cleanup0Base`
The `-cleanup0Base` flag of the deinstallation tool is deprecated in this release. There is no replacement for this flag.

Desupported Features

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

- Oracle Enterprise Manager Database Control
- CLEANUP_ORACLE_BASE Property Removed

Other Changes

- Document Structure Changes
This book is redesigned to provide an installation checklist for Oracle Grid Infrastructure installation, which comprises Oracle Clusterware and Oracle Automatic Storage Management installation. Use the checklist to prepare for installation. For more details, refer to the chapters that subdivide preinstallation tasks into category topics.
- Preinstallation Task Changes
To facilitate cluster deployment, Oracle Universal Installer (OUI) and Cluster Verification Utility (CVU) detects when minimum requirements for installation are not completed, and creates shell script programs, called Fixup scripts, to resolve many incomplete system configuration requirements. If OUI detects an incomplete task that is marked "fixable", then you can easily fix the issue by clicking Fix & Check Again to generate a Fixup script.

Fixup scripts do not replace system tuning, but they do reduce the amount of manual system configuration required for an initial deployment. For this reason, some manual tasks that Fixup scripts perform are now moved to an appendix. If you choose to, you can continue to configure your servers manually.

See [Section 4.4, "Using Installation Fixup Scripts"](#) and [Appendix F, "How to Complete Preinstallation Tasks Manually"](#)

- **Desupport of 32-bit Platforms**

Oracle Grid Infrastructure and Oracle Real Application Clusters can no longer be installed on 32-bit systems.

Oracle Grid Infrastructure Installation Checklist

The following checklist provides a list of required preinstallation steps.

Use this checklist to coordinate tasks to help to ensure that all system and storage preparation and configuration tasks are completed before starting Oracle Grid Infrastructure for a cluster installation.

1.1 System Hardware, Software and Configuration Checklists

This section contains the following server configuration installation checklists:

- [Oracle Grid Infrastructure Installation Server Hardware Checklist](#)
- [Oracle Grid Infrastructure and Oracle RAC Environment Checklist](#)
- [Oracle Grid Infrastructure Network Checklist](#)
- [Oracle Grid Infrastructure and Oracle RAC Upgrades Checklist](#)
- [Oracle Grid Infrastructure Storage Configuration Tasks](#)
- [Oracle Grid Infrastructure Starting the Installation Tasks](#)

1.1.1 Oracle Grid Infrastructure Installation Server Hardware Checklist

Review the following hardware checklist for all installations:

Table 1–1 Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Server hardware	Server make, model, core architecture, and host bus adaptors (HBA) are supported to run with Oracle RAC.
Network Switches	<ul style="list-style-type: none"> ▪ Public network switch, at least 1 GbE, connected to a public gateway. ▪ Private network switch, at least 1 GbE, with 10 GbE recommended, dedicated for use only with other cluster member nodes. The interface must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP. Alternatively, use InfiniBand for the interconnect.
Runlevel	Servers should be either in runlevel 3 or runlevel 5.
Random Access Memory (RAM)	At least 4 GB of RAM for Oracle Grid Infrastructure for a Cluster installations, including installations where you plan to install Oracle RAC.
Temporary disk space allocation	At least 1 GB allocated to /tmp.

Table 1–1 (Cont.) Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Operating System	<ul style="list-style-type: none"> ▪ Supported in the list of supported kernels and releases listed in "About Operating System Requirements" on page 4-9. ▪ Same operating system kernel running on each cluster member node. ▪ OpenSSH installed manually, if you do not have it installed already as part of a default Linux installation, as described in "Open SSH Requirement for Minimal Installation".
Storage hardware	Either Storage Area Network (SAN) or Network-Attached Storage (NAS).
Local Storage Space for Oracle Software	<ul style="list-style-type: none"> ▪ At least 8 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 12 GB of space for the Oracle base of the Oracle Grid Infrastructure installation owner (Grid user). The Oracle base includes Oracle Clusterware and Oracle ASM log files. ▪ 10 GB of additional space in the Oracle base directory of the Grid Infrastructure owner for diagnostic collections generated by Trace File Analyzer (TFA) Collector. ▪ For Linux platforms, if you intend to install Oracle Database, then allocate 6.4 GB of disk space for the Oracle home (the location for the Oracle Database software binaries).
Intelligent Platform Management Interface (IPMI)	<p>Configuration completed, with IPMI administrator account information available to the person running the installation.</p> <p>If you intend to use IPMI, then ensure baseboard management controller (BMC) interfaces are configured, and have an administration account username and password to provide when prompted during installation.</p> <p>For nonstandard installations, if you must change configuration on one or more nodes after installation (for example, if you have different administrator user names and passwords for BMC interfaces on cluster nodes), then decide if you want to reconfigure the BMC interface, or modify IPMI administrator account information after installation.</p>

1.1.2 Oracle Grid Infrastructure and Oracle RAC Environment Checklist

Review the following environment checklist for all installations:

Table 1–2 Environment Configuration for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Create Groups and Users	Review Section 6.1, "Creating Groups, Users and Paths for Oracle Grid Infrastructure" for information about the groups and users you need to create for the kind of deployment you want to do. Installation owners have resource limits settings and other requirements. Group and user names must use only ASCII characters.
Create mount point paths for the software binaries	Oracle recommends that you follow the guidelines for an Optimal Flexible Architecture configuration, as described in the appendix "Optimal Flexible Architecture," in <i>Oracle Database Installation Guide</i> for your platform.
Review Oracle Inventory (oraInventory) and OINSTALL Group Requirements	<p>The Oracle Inventory directory is the central inventory of Oracle software installed on your system. Users who have the Oracle Inventory group as their primary group are granted the OINSTALL privilege to 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, and your system does not have an oraInventory directory, then the installer creates an Oracle inventory that is one directory level up from the Oracle base for the Oracle Grid Infrastructure install, and designates the installation owner's primary group as the Oracle Inventory group. Ensure that this group is available as a primary group for all planned Oracle software installation owners.

Table 1–2 (Cont.) Environment Configuration for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Grid Home Path	<p>Ensure that the Grid home (the Oracle home path you select for Oracle Grid Infrastructure) uses only ASCII characters</p> <p>This 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.</p>
Unset Oracle software environment variables	<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.</p> <p>If you have had an existing installation on your system, and you are using the same user account to install this installation, then unset the following environment variables: <code>ORA_CRS_HOME</code>; <code>ORACLE_HOME</code>; <code>ORA_NLS10</code>; <code>TNS_ADMIN</code>.</p>
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 for all cluster member nodes must use the same password.</p> <ul style="list-style-type: none"> ▪ Use root user credentials <p>Provide the superuser password for cluster member node servers.</p> ▪ Use Sudo <p>Sudo is a UNIX and Linux utility that allows members of the sudoers list privileges to run individual commands as root. Provide the username and password of an operating system user that is a member of sudoers, and is authorized to run Sudo on each cluster member node.</p> <p>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 username and password when prompted during installation.</p> ▪ Run root scripts manually <p>If you run root scripts manually, then you must run the <code>root.sh</code> script on the first node and wait for it to finish. You can then run <code>root.sh</code> concurrently on all other nodes.</p>

1.1.3 Oracle Grid Infrastructure Network Checklist

Review this network checklist for all installations 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 1–3 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 TCP/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. 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). ■ The switches and network interface adapters must be at least 1 GbE, with 10 GbE recommended. Alternatively, use InfiniBand for the interconnect. ■ The interconnect must support the user datagram protocol (UDP).
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. ASM networks use the TCP protocol.
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 on the 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. SCANS 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>
Standard or Hub Node Public, Private and Virtual IP names and Addresses	<p>If you are not using GNS, and you are configuring a Standard cluster, then configure the following for each Hub Node:</p> <ul style="list-style-type: none"> ■ Public node name and address, configured on the DNS and in <code>/etc/hosts</code> (for example, <code>node1.example.com</code>, address <code>192.0.2.10</code>). 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> ■ Public node virtual IP name and address (for example, <code>node1-vip.example.com</code>, address <code>192.0.2.11</code>). <p>If you are not using GNS, 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>

1.1.4 Oracle Grid Infrastructure and Oracle RAC Upgrades Checklist

Review this upgrade checklist if you have an existing Oracle Grid Infrastructure or Oracle RAC installation. A cluster is being upgraded until all cluster member nodes

are running Oracle Grid Infrastructure 12c Release 1 (12.1), and the new clusterware becomes the active version.

If you intend to install Oracle RAC, then you must first complete the upgrade to Oracle Grid Infrastructure 12c Release 1 (12.1) on all cluster nodes before you install the Oracle Database 12c Release 1 (12.1) version of Oracle RAC.

Note: All Oracle Grid Infrastructure upgrades (upgrades of existing Oracle Clusterware and Oracle ASM installations) are out-of-place upgrades. You cannot upgrade from an existing Oracle Grid Infrastructure installation to an Oracle Flex Cluster installation.

Table 1–4 Installation Upgrades Checklist for Oracle Grid infrastructure

Check	Task
Read documentation	Review <i>Oracle Database Upgrade Guide</i>
Latest patchset	Install the latest available patchset release.
Installation owner	<p>Confirm that the installation owner you plan to use is the same as the installation owner that owns the existing Oracle Grid Infrastructure installation.</p> <p>The new Oracle Grid Infrastructure installation and the Oracle Grid Infrastructure home installation that you are upgrading must be owned by same operating system user, or permission errors result.</p>
Oracle Automatic Storage Management (Oracle ASM) instances	<p>Confirm that the Oracle Automatic Storage Management (Oracle ASM) instances you have use standard Oracle ASM instance names.</p> <p>The default ASM SID for a single-instance database is +ASM, and the default SID for ASM on Oracle Real Application Clusters nodes is +ASMnode#, where <i>node#</i> is the node number. With Oracle Grid Infrastructure 11.2.0.1 and later, non-default Oracle ASM instance names are not supported.</p> <p>If you have non-default Oracle ASM instance names, then before you upgrade your cluster, use your existing release <code>srvctl</code> to remove individual Oracle ASM instances with non-default names, and add Oracle ASM instances with default names.</p>
Network Addresses for Standard Oracle Grid Infrastructure deployments	<p>Ensure the following about IP addresses for the public and private networks:</p> <ul style="list-style-type: none"> ■ The private and public IP addresses are 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.*.*).
Oracle Cluster Registry (OCR) files	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.
Operating System configuration	Confirm that you are using a supported operating system, kernel release, and all required operating system packages for the new Oracle Grid Infrastructure installation.
Oracle Cluster Registry (OCR) file integrity	Run the <code>ocrcheck</code> command to confirm Oracle Cluster Registry (OCR) file integrity. If this check fails, then repair the OCRs before proceeding.
Oracle 12c Upgrade Companion	<p>Review Oracle 12c Upgrade Companion (My Oracle Support Note 1462240.1) for the most current information regarding other upgrade issues:</p> <p>https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1462240.1</p>
Run the Oracle Database Pre-Upgrade utility	<p>Run this SQL script located in the path <code>\$ORACLE_HOME/rdbms/admin</code>, after you complete Oracle Grid Infrastructure installation to prepare your databases for upgrades.</p> <p>For more information, review My Oracle Support Note 884522.1:</p> <p>https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=884522.1</p>
Run the ORAchk Upgrade Readiness Assessment	<p>Run the ORAchk Upgrade Readiness Assessment tool to obtain automated upgrade-specific health checks for Oracle Grid Infrastructure upgrade.</p> <p>For more information, review My Oracle Support Note 1457357.1:</p> <p>https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1457357.1</p>

1.1.5 Oracle Grid Infrastructure Storage Configuration Tasks

Review the following storage configuration task checklist for all installations:

Table 1–5 Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
Provide paths for Oracle Clusterware files	<p>During installation, you are asked to provide paths for the following 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, either on Oracle ASM (preferred), or on a cluster file system, because the files created during installation must be available to all cluster member nodes.</p> <ul style="list-style-type: none"> ▪ Voting files are files that Oracle Clusterware uses to verify cluster node membership and status. <p>The location for voting files must be owned by the user performing the installation (<code>oracle</code> or <code>grid</code>), and must have permissions set to <code>640</code>.</p> ▪ Oracle Cluster Registry files (OCR) contain cluster and database configuration information for Oracle Clusterware. <p>Before installation, the location for OCR files must be owned by the user performing the installation (<code>grid</code> or <code>oracle</code>). That installation user must have <code>oinstall</code> as its primary group. During installation, the installer creates the OCR files and changes ownership of the path and OCR files to <code>root</code>.</p>

1.1.6 Oracle Grid Infrastructure Starting the Installation Tasks

Table 1–6 Oracle Grid Infrastructure Checks to Perform Before Starting the Installer

Check	Task
Check running Oracle processes, and shut down 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.
Ensure cron jobs do not run during installation	<p>If the installer is running when daily <code>cron</code> jobs start, then you may encounter unexplained installation problems if your <code>cron</code> job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily <code>cron</code> jobs are run, or disable daily <code>cron</code> jobs that perform cleanup until after the installation is completed.</p>
Decide if you want to install other languages	<p>During installation, you are asked if you want translation of user interface text into languages other than the default, which is English. If the language set for the operating system is not supported by the installer, then by default the installer runs in the English language.</p> <p>See <i>Oracle Database Globalization Support Guide</i> for detailed information about character sets and language configuration.</p>

Configuring Servers for Oracle Grid Infrastructure and Oracle RAC

This chapter describes the operating system tasks you must complete on your servers before you install Oracle Grid Infrastructure for a cluster and Oracle Real Application Clusters (Oracle RAC). The values provided in this chapter are installation minimum only. Oracle recommends that you configure production systems in accordance with planned system loads.

This chapter contains the following topics:

- [Checking Server Hardware and Memory Configuration](#)
- [General Server Minimum Requirements](#)
- [Server Storage Minimum Requirements](#)
- [Server Memory Minimum Requirements](#)

2.1 Checking Server Hardware and Memory Configuration

Run the following commands to gather your current system information:

1. To determine the physical RAM size, enter the following command:

```
# 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. To determine the size of the configured swap space, enter the following command:

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

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

3. To determine the amount of space available in the `/tmp` directory, enter the following command:

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

4. To determine the amount of free RAM and disk swap space on the system, enter the following command:

```
# free
```

5. To determine if the system architecture can run the software, enter the following command:

```
# 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 using the following command:

```
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. See [Section 2.4, "Server Memory Minimum Requirements"](#) for more information about shared memory planning.

2.2 General Server Minimum Requirements

- Select servers with the same instruction set architecture; running 32-bit and 64-bit Oracle software versions in the same cluster stack is not supported.
- Ensure that the server is started with runlevel 3 or 5.
- Ensure display cards provide at least 1024 x 768 display resolution, so that OUI displays correctly while performing a system console-based installation.
- Ensure servers run the same operating system binary.
- Oracle Grid Infrastructure installations and Oracle Real Application Clusters (Oracle RAC) support servers with different hardware in the same cluster. Your cluster can have nodes with CPUs of different speeds or sizes, but Oracle recommends that you use nodes with the same hardware configuration.

Oracle recommends that if you configure clusters using different configuration, that you categorize cluster nodes into homogenous pools as part of your server categorization management policy.

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about server state and configuration attributes, and about using server pools to manage resources and workloads

2.3 Server Storage Minimum Requirements

Each system must meet the following minimum storage requirements:

- 1 GB of space in the `/tmp` directory.

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 make available the space required.
- Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.

- At least 8 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.
- For IBM: Linux on System z, at least 3.8 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home).
- At least 12 GB of space for the Oracle base of the Oracle Grid Infrastructure installation owner (Grid user). The Oracle base includes Oracle Clusterware and Oracle ASM log files.
- For IBM: Linux on System z, at least 200 MB of space for the Oracle base of the Oracle Grid Infrastructure installation owner (Grid user).
- 10 GB of additional space in the Oracle base directory of the Grid Infrastructure owner for diagnostic collections generated by Trace File Analyzer (TFA) Collector.
- For Linux x86-64 platforms, if you intend to install Oracle Database, then allocate 6.4 GB of disk space for the Oracle home (the location for the Oracle Database software binaries).
- For IBM: Linux on System z, if you intend to install Oracle Database, then allocate 5.2 GB of disk space for the Oracle home (the location for the Oracle Database software binaries).

If you are installing Oracle Databases, and you plan to configure automated database backups, then you require additional space either in a file system or in an Oracle Automatic Storage Management disk group for the Fast Recovery Area.

See Also: *Oracle Database Backup and Recovery User's Guide* for more information about Fast Recovery Area sizing

2.4 Server Memory Minimum Requirements

Ensure that your system meets the following minimum requirements, depending on your system architecture

- [Minimum Memory Requirements for Oracle Grid Infrastructure](#)
- [Shared Memory Requirements](#)

Note: If you encounter an OUI error indicating inadequate swap space size, but your swap space meets the requirements listed here, then you can ignore that error.

2.4.1 Minimum Memory Requirements for Oracle Grid Infrastructure

Each system must meet the following minimum memory requirements:

- At least 4 GB of RAM for Oracle Grid Infrastructure for a Cluster installations, including installations where you plan to install Oracle RAC.
- Swap space equivalent to the multiple of the available RAM, as indicated in the following table:

Table 2–1 *Swap Space Required for 64-bit Linux and Linux on System z*

Available RAM	Swap Space Required
Between 4 GB and 16 GB	Equal to RAM

Table 2–1 (Cont.) Swap Space Required for 64-bit Linux and Linux on System z

Available RAM	Swap Space Required
More than 16 GB	16 GB of RAM

Note: 32-bit systems are no longer supported.

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.

See Also: [Appendix D, "Configuring Large Memory Optimization"](#)

2.4.2 Shared Memory Requirements

If you intend to install Oracle Databases or Oracle RAC databases on the cluster, be aware that the size of the shared memory mount area (`/dev/shm`) on each server must be greater than the system global area (SGA) and the program global area (PGA) of the databases on the servers. Review expected SGA and PGA sizes with database administrators to ensure that you do not have to increase `/dev/shm` after databases are installed on the cluster.

See Also: [Section 4.13, "Checking Shared Memory File System Mount on Linux"](#)

Automatically Configuring Oracle Linux with Oracle Preinstallation RPM

Oracle recommends that you install Oracle Linux 7, Oracle Linux 6, or Oracle Linux 5 and use Oracle RPMs to configure your operating systems for Oracle Grid Infrastructure and Oracle Database installations with Oracle Real Application Clusters (Oracle RAC). For Oracle Linux 7 and Oracle Linux 6, run the Oracle Preinstallation RPM. For Oracle Linux 5, run the Oracle Validated RPM.

This chapter contains the following contents:

- [Overview of Oracle Linux Configuration with Oracle RPMs](#)
- [Installing the Oracle Preinstallation RPM From Unbreakable Linux Network](#)
- [Installing Oracle Linux with Oracle Linux Yum Server Support](#)
- [Installing the Oracle Preinstallation RPM From DVDs or Images](#)
- [Additional Optional Operating System Configuration Tasks](#)
- [Required System Configuration for Oracle Grid Infrastructure](#)

3.1 Overview of Oracle Linux Configuration with Oracle RPMs

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:

<http://linux.oracle.com>

Configuring a server using Oracle Linux and an Oracle 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 or Oracle Validated 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 on each cluster candidate node.

5. Complete network interface configuration for each cluster node candidate.
6. Complete system configuration for shared storage access as required for each standard or Hub 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 From Unbreakable Linux Network

Use the following procedure to subscribe to Oracle Linux channels, and to add the Oracle Linux channel that distributes the Oracle RDBMS Server 12cR1 RPM:

1. Complete a default Oracle Linux workstation installation, or a default Red Hat Enterprise Linux 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**. The System Summary window opens.
6. From the **Available Channels** list, select the Enterprise Linux installation media copy and update patch channels corresponding to your Oracle Linux distribution. For example, if your distribution is Oracle Linux 5 Update 6 for x86_64, then select the following:
 - Oracle Linux 5 Update 6 installation media copy (x86_64)
 - Oracle Linux 5 Update 6 Patch (x86_64)
7. Click **Subscribe**.
8. Start a terminal session and enter the following command as `root`, depending on your platform:

Oracle Linux 6:

```
yum install oracle-rdbms-server-12cR1-preinstall
```

Oracle Linux 5:

```
# yum install oracle-validated
```

You should see output indicating that you have subscribed to the Oracle Linux channel. For example:

```
e15_u5_i386_base  
e15_u5_x86_64_patch
```

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.

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-rdbms-server-12cR1-preinstall.conf` file and other configuration files.

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

Note: Check the RPM log file to review the system configuration changes. For example, on Oracle Linux 5:

```
/var/log/oracle-validated/results/orakernel.log
```

3.3 Installing Oracle Linux with Oracle Linux Yum Server Support

Use the following 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 <http://yum.oracle.com>, using the instructions you can find on the Oracle Linux yum server. For example:

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

Ensure that the `olrelease_latest` file (`ol6_latest` for Oracle Linux 6) is enabled, as 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 `ol6_UEK_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:

```
yum install oracle-rdbms-server-12cR1-preinstall
```

Oracle Linux 5:

```
# yum install oracle-validated
```

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

```
e15_u6_i386_base  
e15_u6_x86_64_patch
```

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.

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

3.4 Installing the Oracle Preinstallation RPM From DVDs or Images

Use the following procedure to install the Oracle Preinstallation RPM from the Oracle Linux distribution:

1. Obtain Oracle Linux disks 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:

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

2. Start the Oracle Linux installation.
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 7 and Oracle Linux 6, select **Servers** on the left hand side of the screen, and then select **System administration Tools** on the right side of the screen. These options may differ between releases.

The Packages in System Tools window opens.

5. Select the Oracle Preinstallation RPM package box from the package list, and click **Next**.
6. Complete the other screens to finish installing Oracle Linux.

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.

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

3.5 Additional Optional Operating System Configuration Tasks

Complete the following optional configuration tasks:

- [Configure Oracle Ksplice Repository for Oracle Linux](#)
- [Configure Additional Operating System Features](#)

3.5.1 Configure Oracle Ksplice Repository for Oracle Linux

You can use Oracle Ksplice if you have Premier support subscription and an access key, which is available on ULN. For more information about Ksplice (including trial versions, see <http://www.ksplice.com/>).

Complete the following task to register your system with Ksplice:

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

<http://www.ksplice.com/uptrack/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.ksplice.com/yum/uptrack/ol/ksplice-uptrack-release.noarch.rpm>

5. Run the following commands:

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

3.5.2 Configure Additional Operating System Features

As needed, configure the operating system for additional features, such as Intelligent Platform Management Interface (IPMI), or additional programming environments, then review [Chapter 4, "Configuring Operating Systems for Oracle Grid Infrastructure and Oracle RAC."](#)

3.6 Required System Configuration for Oracle Grid Infrastructure

Complete system configuration as described in the following chapters:

- [Chapter 5, "Configuring Networks for Oracle Grid Infrastructure and Oracle RAC"](#)
- [Chapter 6, "Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle RAC"](#)
- [Chapter 7, "Configuring Storage for Oracle Grid Infrastructure and Oracle RAC"](#)

Configuring Operating Systems for Oracle Grid Infrastructure and Oracle RAC

This chapter describes the operating system configuration tasks you must complete on your servers before you install Oracle Grid Infrastructure for a cluster and Oracle Real Application Clusters.

This chapter contains the following topics:

- [Guidelines for Linux Operating System Installation](#)
- [Reviewing Operating System and Software Upgrade Best Practices](#)
- [Reviewing Operating System Security Common Practices](#)
- [Using Installation Fixup Scripts](#)
- [Logging In to a Remote System Using X Terminal](#)
- [Using Oracle RPM Checker on IBM: Linux on System z](#)
- [About Operating System Requirements](#)
- [Operating System Requirements for x86-64 Linux Platforms](#)
- [Operating System Requirements for IBM: Linux on System z](#)
- [Additional Drivers and Software Packages for Linux](#)
- [Checking the Software Requirements](#)
- [Installing the cvuqdisk RPM for Linux](#)
- [Checking Shared Memory File System Mount on Linux](#)
- [Enabling the Name Service Cache Daemon](#)
- [Setting the Disk I/O Scheduler on Linux](#)
- [Setting Network Time Protocol for Cluster Time Synchronization](#)
- [Using Automatic SSH Configuration During Installation](#)

4.1 Guidelines for Linux Operating System Installation

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

This section contains the following topics:

- [Completing a Minimal Linux Installation](#)

- [Completing a Default Linux Installation](#)
- [About Unbreakable Enterprise Kernel for Oracle Linux](#)
- [About the Oracle Preinstallation RPM](#)
- [Using Oracle Ksplice to Perform a Zero Downtime Update](#)

4.1.1 Completing a Minimal Linux Installation

Review the following sections regarding minimal Linux installation requirements:

- [About Minimal Linux Installations](#)
- [RPM Packages for Completing Operating System Configuration](#)
- [Open SSH Requirement for Minimal Installation](#)

4.1.1.1 About Minimal Linux Installations

To complete a minimal Linux installation, select one of the minimal installation options (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). This 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/CSP/main/article?cmd=show&type=NOT&id=728346.1>

Note: 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:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=579101.1>

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

Oracle Linux 6 Preinstallation RPM With ULN Support

Oracle Preinstallation RPM 12c Release 1 (12.1) for your Oracle Linux 6 kernel.

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

Oracle Linux 5 Oracle Validated RPM With ULN Support

Oracle Validated RPM (oracle-validated) for your Oracle Linux 5 kernel.

Unbreakable Linux Network (ULN) customers can obtain the Oracle Validated RPM by using `up2date`, or using `yum` (5.5 and later releases).

See Also: [Chapter 3, "Automatically Configuring Oracle Linux with Oracle Preinstallation RPM"](#)

Oracle Linux 6 Preinstallation RPM Without ULN Support

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

Oracle Linux 5 Oracle Validated RPM Without ULN Support

http://yum.oracle.com/repo/OracleLinux/OL5/latest/x86_64/

4.1.1.3 Open SSH Requirement for Minimal Installation

SSH 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). This 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 Unbreakable Enterprise Kernel for Oracle Linux

Unbreakable Enterprise Kernel for Oracle Linux delivers the latest innovations from upstream development to customers who run Oracle Linux in the data center. The Unbreakable Enterprise Kernel for Oracle Linux is included and enabled by default starting with Oracle Linux 5 Update 6.

The Unbreakable Enterprise Kernel for Oracle Linux 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 for Oracle Linux in your Linux environment, especially if you run enterprise applications. However, using Unbreakable Enterprise Kernel for Oracle Linux is optional. If you require strict 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/index.html>

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. Oracle Unbreakable Enterprise Kernel for Linux is used in all benchmark tests on Linux in which Oracle participates, as well as in the Oracle preinstallation RPM 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, 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 it is 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 preinstallation RPM
- 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 and higher, 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.

See Also: [Chapter 3, "Automatically Configuring Oracle Linux with Oracle Preinstallation RPM"](#)

4.1.5 Using Oracle Ksplice to Perform a Zero Downtime Update

Oracle Ksplice Uptrack updates provide Linux security and bug fix updates, repackaged in a form that allows these updates to be applied without restarting the kernel.

To use Ksplice Uptrack:

1. Obtain or verify your Oracle Linux premium support subscription from Unbreakable Linux Network:

<https://linux.oracle.com>

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.ksplice.com/yum/uptrack/ol/ksplice-uptrack-release.noarch.rpm>

5. Run the following commands:

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

6. Open `/etc/uptrack/uptrack.conf` with a text editor, enter your premium support access key, and save the file. You must use the same access key for all of your systems.

7. Run the following command to carry out a zero downtime update of your kernel:

```
uptrack-upgrade -y
```

See Also:

- The Oracle Ksplice Uptrack website for more information:
<http://www.ksplice.com>
- Oracle Ksplice for Oracle Linux:
<http://oss.oracle.com/ksplice/docs/ksplice-quickstart.pdf>

4.2 Reviewing Operating System and Software Upgrade Best Practices

Review the following information regarding upgrades:

- [General Upgrade Best Practices](#)
- [Oracle ASM Upgrade Notifications](#)
- [Rolling Upgrade Procedure Notifications](#)

4.2.1 General Upgrade Best Practices

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

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

- 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

To find the most recent software updates, and to find best practices recommendations about preupgrade, postupgrade, compatibility, and interoperability, see Oracle 12c Upgrade Companion (My Oracle Support Note 1462240.1):

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1462240.1>

See Also: [Appendix B, "How to Upgrade to Oracle Grid Infrastructure 12c Release 1"](#)

4.2.2 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 1 (11.1) 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. If you are upgrading from Oracle ASM 10g, then you must shut down the entire Oracle ASM cluster to perform the upgrade.
- 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).

If you have an existing Oracle ASM home from a previous release, then it should be owned by the same user that you plan to use to upgrade Oracle Clusterware.

4.2.3 Rolling Upgrade Procedure Notifications

Be aware of the following information regarding rolling upgrades:

- During rolling upgrades of the operating system, Oracle supports using different operating system binaries when both versions of the operating system are certified with the Oracle Database release you are using.
- Using mixed operating system versions is supported during upgrade only.

Be aware that mixed operating systems are supported only supported for the duration of an upgrade, over the period of a few hours.

- Oracle Clusterware does not support nodes that have processors with different instruction set architectures (ISAs) in the same cluster. Each node must be binary compatible with the other nodes in the cluster.

For example, you cannot have one node using an Intel 64 processor and another node using an IA-64 (Itanium) processor in the same cluster. You could have one node using an Intel 64 processor and another node using an AMD64 processor in the same cluster because the processors use the same x86-64 ISA and run the same binary version of Oracle software.

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.

See Also: <http://docs.oracle.com/en/operating-systems/>

4.4 Using Installation Fixup Scripts

Oracle Universal Installer (OUI) 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 OUI detects an incomplete task, then it generates a Fixup script (`runfixup.sh`). You can run the script after you click **Fix and Check Again**.

You also can have CVU generate Fixup scripts before installation.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about using the `cluvfy` command

Fixup scripts do the following:

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

If you have SSH configured between cluster member nodes for the user account that you will use for installation, then you can check your cluster configuration before installation and generate a fixup script to make operating system changes before starting the installation.

To do this, log in as the user account that will perform the installation, navigate to the staging area where the `runcluvfy` command is located, and use the following command syntax, where *node* is a comma-delimited list of nodes you want to make cluster members:

```
$ ./runcluvfy.sh stage -pre crsinst -n node -fixup -verbose
```

For example, if you intend to configure a two-node cluster with nodes `node1` and `node2`, enter the following command:

```
$ ./runcluvfy.sh stage -pre crsinst -n node1,node2 -fixup -verbose
```

4.5 Logging In to a Remote System Using X Terminal

During installation, you are required to perform tasks as root or as other users on remote terminals. Complete the following procedure for user accounts that you want to enable for remote display.

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

To enable remote display, complete one of the following procedures:

- If you are installing the software from an X Window System workstation or X terminal, then:
 1. Start an X terminal session (`xterm`).
 2. If you are installing the software on another system and using the system as an X11 display, then enter a command using the following syntax to enable remote hosts to display X applications on the local X server:

```
# xhost + RemoteHost
```

where *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
```

where *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, then enter the following command to switch the user to `root`:

```
$ su - root
```

```
password:
#
```

- If you are installing the software from a PC or other system with X server software installed, then:

Note: If necessary, see your X Window System documentation for more information about completing this procedure. Depending on the X software that you are using, you may need to complete the tasks in a different order.

1. Start the X Window System software.
2. Configure the security settings of the X Window System software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software as the Oracle Grid Infrastructure for a cluster software owner (`grid`, `oracle`) and start a terminal session on that system; for example, an X terminal (`xterm`).
4. Open another terminal on the remote system, and log in as the `root` user on the remote system, so you can run scripts as `root` when prompted.

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 Oracle Grid Infrastructure installation.

Download the Oracle RPM Checker utility from the link in My Oracle Support note 1574412.1 available at the following URL:

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

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

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

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

4.7 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 on the Oracle Technology Network at the following URL:

<http://www.oracle.com/technetwork/indexes/documentation/index.html>

Oracle Universal Installer (OUI) 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.

4.8 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. No other Linux distributions are supported.

Identify operating system requirements for Oracle Grid Infrastructure, and identify additional operating system requirements for Oracle Database and Oracle RAC installations.

- [Supported Oracle Linux 7 and Red Hat Linux 7 Distributions for x86-64](#)
- [Supported Oracle Linux 6 and Red Hat Linux 6 Distributions for x86-64](#)
- [Supported Oracle Linux 5 and Red Hat Linux 5 Distributions for x86-64](#)
- [Supported SUSE Linux Enterprise Server Distributions for x86-64](#)
- [Supported NeoKylin Linux Advanced Server Distributions for x86-64](#)

Note: 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 might 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/>

Note: ■The Unbreakable Enterprise Kernel for Oracle Linux can be installed on x86-64 servers running either Oracle Linux 5 Update 6, or Red Hat Enterprise Linux 5 Update 6. As of Oracle Linux 5 Update 6, the Unbreakable Enterprise Kernel for Oracle Linux is the default system kernel. An x86 (32-bit) release of Oracle Linux including the Unbreakable Enterprise Kernel for Oracle Linux is available with Oracle Linux 5 update 7 and later.

- The 32-bit packages listed in the following sections are required only for 32-bit client installs.
 - Oracle Universal Installer requires an X Window System (for example, `libx`). The `libx` packages are part of a default Linux installation. If you install Linux using an Oracle Preinstallation RPM, then the `libx` packages are installed as part of that RPM. If you perform an install on a system with a reduced set of packages, then you must ensure that `libx` is installed.
-
-

4.8.1 Supported Oracle Linux 7 and Red Hat Linux 7 Distributions for x86-64

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

Note: Starting with Oracle Database 12c Release 1 (12.1.0.2), Oracle Linux 7 and Red Hat Enterprise Linux 7 are supported on Linux x86-64 systems.

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

- My Oracle Support Note 2081410.1:
<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=2081410.1>
- Oracle Linux 7 Update 2 Release Notes:
<http://docs.oracle.com/en/operating-systems/>

Table 4–1 x86-64 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 public yum site, 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 for Oracle Linux: 3.8.13-33.el7uek.x86_64 or later ■ Oracle Linux 7 with the Red Hat Compatible kernel: 3.10.0-123.el7.x86_64 or later
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

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

Item	Requirements
Packages for Oracle Linux 7 and Red Hat Enterprise Linux 7	<p>The following packages (or later versions) must be installed:</p> <pre> 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 gcc-4.8.2-3.el7.x86_64 gcc-c++-4.8.2-3.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 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 ksh make-3.82-19.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 sysstat-10.1.5-1.el7.x86_64 nfs-utils-1.3.0-0.21.el7.x86_64 </pre>

4.8.2 Supported Oracle Linux 6 and Red Hat Linux 6 Distributions for x86-64

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

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

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Oracle Linux 6	<p>Subscribe to the Oracle Linux 6 channel on the Unbreakable Linux Network, or configure a yum repository from the Oracle public yum site, 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>Supported distributions:</p> <ul style="list-style-type: none"> ■ Oracle Linux 6 with the Unbreakable Enterprise Kernel for Oracle Linux: <ul style="list-style-type: none"> Update 2 or higher, 2.6.39-200.24.1.el6uek.x86_64 or later UEK2 kernels Update 4 or higher, 3.8.13-16 or later UEK3 kernels Update 7 or higher, 4.1.12-32 or later UEK4 kernels ■ Oracle Linux 6 with the Red Hat Compatible kernel: 2.6.32-71.el6.x86_64 or later

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

Item	Requirements
Red Hat Enterprise Linux 6	Supported distributions: <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux 6: 2.6.32-71.el6.x86_64 or later
Packages for Oracle Linux 6 and Red Hat Enterprise Linux 6	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.20.51.0.2-5.11.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 gcc-4.4.4-13.el6 (x86_64) gcc-c++-4.4.4-13.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 libXext-1.1 (x86_64) libXext-1.1 (i686) libXtst-1.0.99.2 (x86_64) libXtst-1.0.99.2 (i686) libX11-1.3 (x86_64) libX11-1.3 (i686) libXau-1.0.5 (x86_64) libXau-1.0.5 (i686) libxcb-1.5 (x86_64) libxcb-1.5 (i686) libXi-1.3 (x86_64) libXi-1.3 (i686) make-3.81-19.el6 sysstat-9.0.4-11.el6 (x86_64) nfs-utils-1.2.3-15.0.1 </pre>

4.8.3 Supported Oracle Linux 5 and Red Hat Linux 5 Distributions for x86-64

Use the following information to check supported Oracle Linux 5 and Red Hat Linux 5 distributions:

Table 4–3 x86-64 Linux 5 Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
Oracle Linux 5	<p>Subscribe to the Oracle Linux 5 channel on the Unbreakable Linux Network, and then install the Oracle Validated RPM. This RPM installs all required kernel packages for Oracle Grid Infrastructure and Oracle Database installations, and performs other system configuration.</p> <p>Supported distributions:</p> <ul style="list-style-type: none"> ▪ Oracle Linux 5 Update 6 with the Unbreakable Enterprise Kernel for Oracle Linux: 2.6.32-100.0.19 or later ▪ Oracle Linux 5 Update 6 with the Red Hat compatible Kernel: 2.6.18-238.0.0.0.1.el5
Red Hat Enterprise Linux 5	<p>Supported distributions:</p> <ul style="list-style-type: none"> ▪ Red Hat Enterprise Linux 5 Update 6: 2.6.18-238.0.0.0.1.el5 or later
Package requirements for Oracle Linux 5 and Red Hat Enterprise Linux 5	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.17.50.0.6 compat-libstdc++-33-3.2.3 compat-libstdc++-33-3.2.3 (32 bit) gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-58 glibc-2.5-58 (32 bit) glibc-devel-2.5-58 glibc-devel-2.5-58 (32 bit) ksh libaio-0.3.106 libaio-0.3.106 (32 bit) libaio-devel-0.3.106 libaio-devel-0.3.106 (32 bit) libgcc-4.1.2 libgcc-4.1.2 (32 bit) libstdc++-4.1.2 libstdc++-4.1.2 (32 bit) libstdc++-devel 4.1.2 libXext-1.0.1 libXext-1.0.1 (32 bit) libXtst-1.0.1 libXtst-1.0.1 (32 bit) libX11-1.0.3 libX11-1.0.3 (32 bit) libXau-1.0.1 libXau-1.0.1 (32 bit) libXi-1.0.1 libXi-1.0.1 (32 bit) make-3.81 sysstat-7.0.2 nfs-utils-1.0.9-60.0.2 coreutils-5.97-23.el5_4.1 </pre>

4.8.4 Supported SUSE Linux Enterprise Server Distributions for x86-64

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

Table 4–4 x86-64 Supported SUSE Linux Enterprise Server 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: <ul style="list-style-type: none"> ■ SUSE Linux Enterprise Server 12 SP1: 3.12.49-11 or later ■ SUSE Linux Enterprise Server 11 SP2: 3.0.13-0.27 or later
SUSE Linux Enterprise Server 12	The following packages (or later versions) must be installed: binutils-2.25.0-13.1 gcc-4.8-6.189 gcc48-4.8.5-24.1 glibc-2.19-31.9 glibc-32bit-2.19-31.9 glibc-devel-2.19-31.9.x86_64 glibc-devel-32bit-2.19-31.9.x86_64 mksh-50-2.13 libaio1-0.3.109-17.15 libaio-devel-0.3.109-17.15 libcap1-1.10-59.61 libstdc++48-devel-4.8.5-24.1.x86_64 libstdc++48-devel-32bit-4.8.5-24.1.x86_64 libstdc++6-5.2.1+r226025-4.1.x86_64 libstdc++6-32bit-5.2.1+r226025-4.1.x86_64 libstdc++-devel-4.8-6.189.x86_64 libstdc++-devel-32bit-4.8-6.189.x86_64 libgcc_s1-5.2.1+r226025-4.1.x86_64 libgcc_s1-32bit-5.2.1+r226025-4.1.x86_64 make-4.0-4.1.x86_64 sysstat-10.2.1-3.1.x86_64 xorg-x11-driver-video-7.6_1-14.30.x86_64 xorg-x11-server-7.6_1.15.2-36.21.x86_64 xorg-x11-essentials-7.6_1-14.17.noarch xorg-x11-Xvnc-1.4.3-7.2.x86_64 xorg-x11-fonts-core-7.6-29.45.noarch xorg-x11-7.6_1-14.17.noarch xorg-x11-server-extra-7.6_1.15.2-36.21.x86_64 xorg-x11-libs-7.6-45.14.noarch xorg-x11-fonts-7.6-29.45.noarch

Table 4–4 (Cont.) x86-64 Supported SUSE Linux Enterprise Server Operating System Requirements

Item	Requirements
SUSE Linux Enterprise Server 11	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.21.1-0.7.25 gcc-4.3-62.198 gcc-c++-4.3-62.198 glibc-2.11.3-17.31.1 glibc-devel-2.11.3-17.31.1 ksh-93u-0.6.1 libaio-0.3.109-0.1.46 libaio-devel-0.3.109-0.1.46 libcap1-1.10-6.10 libstdc++33-3.3.3-11.9 libstdc++33-32bit-3.3.3-11.9 libstdc++43-devel-4.3.4_20091019-0.22.17 libstdc++46-4.6.1_20110701-0.13.9 libgcc46-4.6.1_20110701-0.13.9 make-3.81 sysstat-8.1.5-7.32.1 xorg-x11-libs-32bit-7.4 xorg-x11-libs-7.4 xorg-x11-libX11-32bit-7.4 xorg-x11-libX11-7.4 xorg-x11-libXau-32bit-7.4 xorg-x11-libXau-7.4 xorg-x11-libxcb-32bit-7.4 xorg-x11-libxcb-7.4 xorg-x11-libXext-32bit-7.4 xorg-x11-libXext-7.4 nfs-kernel-server-1.2.1-2.24.1.x86_64 </pre>

4.8.5 Supported NeoKylin Linux Advanced Server Distributions for x86-64

Use the following information to check supported NeoKylin Linux Advanced Server distributions:

Table 4–5 x86-64 Supported NeoKylin Linux Minimum Operating System Requirements

Item	Requirements
SSH Requirement	Ensure that OpenSSH is installed on your servers. OpenSSH is the required SSH software.
NeoKylin Linux Advanced Server	<p>Supported distributions:</p> <p>NeoKylin Linux Advanced Server 6: 2.6.32-431.el6.x86_64 or later</p>

Table 4–5 (Cont.) x86-64 Supported NeoKylin Linux Minimum Operating System Requirements

Item	Requirements
NeoKylin 6.0	<p>The following packages (or later versions) must be installed:</p> <pre> 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) gcc-4.4.7-4.el6 (x86_64) gcc-c++-4.4.7-4.el6 (x86_64) glibc-2.12-1.132.el6 (i686) glibc-2.12-1.132.el6 (x86_64) glibc-devel-2.12-1.132.el6 (x86_64) glibc-devel-2.12-1.132.el6 (i686) ksh libgcc-4.4.7-4.el6 (i686) libgcc-4.4.7-4.el6 (x86_64) libstdc++-4.4.7-4.el6 (x86_64) libstdc++-4.4.7-4.el6 (i686) libstdc++-devel-4.4.7-4.el6 (x86_64) libstdc++-devel-4.4.7-4.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) libXext-1.3.1-2.el6 (x86_64) libXext-1.3.1-2.el6 (i686) libXtst-1.2.1-2.el6 (x86_64) libXtst-1.2.1-2.el6 (i686) libX11-1.5.0-4.el6 (x86_64) libX11-1.5.0-4.el6 (i686) libXau-1.0.6-4.el6 (x86_64) libXau-1.0.6-4.el6 (i686) libxcb-1.8.1-1.el6 (x86_64) libxcb-1.8.1-1.el6 (i686) libXi-1.6.1-3.el6 (x86_64) libXi-1.6.1-3.el6 (i686) make-3.81-20.el6 sysstat-9.0.4-22.el6 (x86_64) </pre>

4.9 Operating System Requirements for IBM: Linux on System z

The distributions and packages listed in this section are supported for this release on IBM: Linux on System z. No other IBM: Linux on System z distributions are supported.

Identify operating system requirements for Oracle Grid Infrastructure, and identify additional operating system requirements for Oracle Database and Oracle RAC installations.

- [Supported Red Hat Enterprise Linux 7 Distributions for IBM: Linux on System z](#)
- [Supported Red Hat Enterprise Linux 6 Distributions for IBM: Linux on System z](#)
- [Supported Red Hat Enterprise Linux 5 Distributions for IBM: Linux on System z](#)
- [Supported SUSE Distributions for IBM: Linux on System z](#)

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

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

Table 4–6 IBM: Linux on System z 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	<p>Red Hat Enterprise Linux 7: 3.10.0-229.el7.s390x or later</p> <p>Note: You can install on Red Hat Enterprise Linux 7 Update 1, but Oracle recommends that you install on Red Hat Enterprise Linux 7 Update 2 for seamless security enhancements.</p> <p>See My Oracle Support Note 2213265.1 for more information: https://support.oracle.com/rs?type=doc&id=2213265.1</p>
Packages for Red Hat Enterprise Linux 7	<p>The following packages (or later versions) must be installed:</p> <pre>binutils-2.23.52.0.1-30.el7.s390x compat-libcap1-1.10-7.el7.s390x compat-libstdc++-33-3.2.3-71.el7 (s390) compat-libcap1-1.10-1 (s390x) cpp-4.8.2-16.el7.s390x gcc-4.8.3-9.el7.s390x gcc-4.8.3-9.el7.s390x glibc-2.17-78.el7 (s390) glibc-devel-2.17-78.el7 (s390x) glibc-devel-2.17-78.el7 (s390) glibc-devel-2.17-78.el7 (s390x) glibc-headers-2.17-78.el7 (s390x) ksh-20120801-22.el7 (s390x) libaio-0.3.109-12.el7 (s390) libaio-0.3.109-12.el7 (s390) libaio-devel-0.3.109-12.el7 (s390x) libgcc-4.8.3-9.el7 (s390) libgcc-4.8.3-9.el7 (s390x) libstdc++-4.8.3-9.el7 (s390) libstdc++-4.8.3-9.el7 (s390x) libstdc++-devel-4.8.3-9.el7 (s390) libstdc++-devel-4.8.3-9.el7 (s390x) libXtst-1.2.2-2.1.el7 (s390) libXtst-1.2.2-2.1.el7 (s390x) libXi-1.7.2-2.1.el7 (s390x) libXi-1.7.2-2.1.el7 (s390x) libxcb-1.9-5.1.el7 (s390) libxcb-1.9-5.el7 (s390x) llibX11-1.6.0-2.el7 (s390) libX11-1.6.0-2.el7 (s390x) libXau-1.0.8-2.1.el7 (s390) libXau-1.0.8-2.1.el7 (s390x) libXext-1.3.2-2.1.el7 (s390) libXext-1.3.2-2.1.el7 (s390x) make-3.82-21.el7 (s390x) mpfr-3.1.1-4.el7.s390x sysstat-10.1.5-7.el7 (s390x)</pre>

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

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

Table 4–7 IBM: Linux on System z 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	<p>Red Hat Enterprise Linux 6: 2.6.32-279.el6.s390x or later</p> <p>Note: You can install on Red Hat Enterprise Linux 6 Update 3, but Oracle recommends that you install on Red Hat Enterprise Linux 6 Update 4 as RHEL 6.4 includes significant I/O performance gains on Open Storage.</p> <p>See My Oracle Support Note 1574412.1 for more information: https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=2081410.1</p>
Packages for Red Hat Enterprise Linux 6	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.20.51.0.2-5.28 (s390x) compat-libstdc++-33-3.2.3-69.el6 (s390x) compat-libcap1-1.10-1 (s390x) gcc-4.4.6-3.el6 (s390x) gcc-c++-4.4.6-3.el6 (s390x) glibc-2.12-1.80.el6 (s390) glibc-2.12-1.80.el6 (s390x) glibc-devel-2.12-1.80.el6 (s390) glibc-devel-2.12-1.80.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-4.4.6-4.el6 (s390) libgcc-4.4.6-4.el6 (s390x) libstdc++-4.4.6-4.el6 (s390x) libstdc++-devel-4.4.6-4.el6 (s390x) libXtst-1.0.99.2-3.el6 (s390) libXtst-1.0.99.2-3.el6 (s390x) libXi-1.3-3.el6 (s390) libXi-1.3-3.el6 (s390x) libXmu-1.0.5-1.el6 (s390) libXaw-1.0.6-4.1.el6 (s390) libXft-2.1.13-4.1.el6 (s390) libXp-1.0.0-15.1.el6 (s390) make-3.81-20.el6 (s390x) ksh-20100621-16.el6 (s390x) sysstat-9.0.4-18.el6 (s390x) </pre>

4.9.3 Supported Red Hat Enterprise Linux 5 Distributions for IBM: Linux on System z

Use the following information to check supported Red Hat Linux 5 distributions:

Table 4–8 IBM: Linux on System z Linux 5 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 5	Red Hat Enterprise Linux 5.8: 2.6.18-308.el5 s390x or later
Package requirements for Red Hat Enterprise Linux 5	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.17.50.0.6-20.el5 (s390x) compat-libstdc++-33-3.2.3-61 (s390) compat-libstdc++-33-3.2.3-61 (s390x) gcc-c++-4.1.2-52.el5 (s390x) gcc44-4.4.6-3.el5.1 (s390x) glibc-2.5-81 (s390) glibc-2.5-81 (s390x) glibc-devel-2.5-81 (s390) glibc-devel-2.5-81 (s390x) libaio-0.3.106-5 (s390) libaio-0.3.106-5 (s390x) libaio-devel-0.3.106-5 (s390) libaio-devel-0.3.106-5 (s390x) libgcc-4.1.2-52.el5 (s390) libgcc-4.1.2-52.el5 (s390x) libstdc++-4.1.2-52.el5 (s390) libstdc++-4.1.2-52.el5 (s390x) libstdc++-devel-4.1.2-52.el5 (s390x) libstdc++44-devel-4.4.6-3.el5.1 (s390) libstdc++44-devel-4.4.6-3.el5.1 (s390x) libXi-1.0.1-4.el5_4 (s390) libXi-1.0.1-4.el5_4 (s390x) libXtst-1.0.1-3.1 (s390) libXtst-1.0.1-3.1 (s390x) make-3.81-3.el5 (s390x) ksh-20100621-5.el5 (s390x) sysstat-7.0.2-11.el5 (s390x) </pre>

4.9.4 Supported SUSE Distributions for IBM: Linux on System z

Use the following information to check supported SUSE distributions:

Table 4–9 IBM: Linux on System z SUSE 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	SUSE Linux Enterprise Server 12 SP1: 3.12.53-60.30.1-default s390x or later SUSE Linux Enterprise Server 11 SP2: 3.0.13-0.27-default s390x or later
SUSE 12	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.25.0-13.1 (s390x) gcc-32bit-4.8-6.189 (s390x) gcc48-4.8.5-24.1 (s390x) gcc48-32bit-4.8.5-24.1 (s390x) gcc48-c++-4.8.5-24.1 (s390x) gcc48-info-4.8.5-24.1 (s390x) gcc48-locale-4.8.5-24.1 (s390x) gcc-c++-4.8-6.189 (s390x) gcc-c++-32bit-4.8-6.189 (s390x) gcc-info-4.8-6.189 (s390x) gcc-locale-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) libaio1-0.3.109-17.15 (s390x) libaio1-32bit-0.3.109-17.15 (s390x) libaio-devel-0.3.109-17.15 (s390x) libcap1-1.10-59.61 (s390x) libcap1-32bit-1.10-59.61 (s390x) libcap2-2.22-11.709 (s390x) libcap2-32bit-2.22-11.709 (s390x) libcap-ng0-0.7.3-4.125 (s390x) libcap-ng0-32bit-0.7.3-4.125 (s390x) libcap-ng-utils-0.7.3-4.125 (s390x) libcap-progs-2.22-11.709 (s390x) libgcc_s1-5.2.1+r226025-4.1 (s390x) libgcc_s1-32bit-5.2.1+r226025-4.1 (s390x) libgomp1-32bit-5.2.1+r226025-4.1 (s390x) libstdc++48-devel-4.8.5-24.1 (s390x) libstdc++48-devel-32bit-4.8.5-24.1 (s390x) libstdc++6-5.2.1+r226025-4.1 (s390x) libstdc++6-32bit-5.2.1+r226025-4.1 (s390x) libstdc++-devel-4.8-6.189 (s390x) libstdc++-devel-32bit-4.8-6.189 (s390x) libXtst6-1.2.2-3.60 (s390x) libXtst6-32bit-1.2.2-3.60 (s390x) make-4.0-4.1 (s390x) mksh-50-2.13 (s390x) sysstat-10.2.1-3.1 (s390x) xorg-x11-7.6_1-14.17 (s390x) xorg-x11-essentials-7.6_1-14.17 (s390x) xorg-x11-fonts-7.6-29.45 (s390x) xorg-x11-fonts-core-7.6-29.45 (s390x) xorg-x11-libs-7.6-45.14 (s390x) xorg-x11-server-7.6_1.15.2-36.21 (s390x) xorg-x11-server-extra-7.6_1.15.2-36.21 (s390x) xorg-x11-Xvnc-1.4.3-7.2 (s390x) OCFS2 1.4 (For Oracle RAC only) </pre>

Table 4–9 (Cont.) IBM: Linux on System z SUSE Minimum Operating System Requirements

Item	Requirements
SUSE 11	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.21.1-0.7.25 (s390x) gcc-4.3-62.198 (s390x) gcc-c++-4.3-62.198 (s390x) glibc-2.11.3-17.31.1 (s390x) glibc-32bit-2.11.3-17.31.1 (s390x) glibc-devel-2.11.3-17.31.1 (s390x) glibc-devel-32bit-2.11.3-17.31.1 (s390x) ksh-93u-0.6.1 (s390x) make-3.81-128.20 libaio-0.3.109-0.1.46 (s390x) libaio-32bit-0.3.109-0.1.46 (s390x) libaio-devel-0.3.109-0.1.46 (s390x) libaio-devel-32bit-0.3.109-0.1.46 (s390x) libcap1-1.10-6.10 (s390x) libgcc46-4.6.1_20110701-0.13.9 (s390x) libstdc++33-3.3.3-11.9 (s390x) libstdc++33-32bit-3.3.3-11.9 (s390x) libstdc++43-devel-32bit-4.3.4_20091019-0.22.17 (s390x) libstdc++43-devel-4.3.4_20091019-0.22.17 (s390x) libstdc++46-32bit-4.6.1_20110701-0.13.9 (s390x) libstdc++46-4.6.1_20110701-0.13.9 (s390x) sysstat-8.1.5-7.32.1 (s390x) xorg-x11-libX11-32bit-7.4-5.9.1 (s390x) xorg-x11-libX11-7.4-5.9.1 (s390x) xorg-x11-libXau-32bit-7.4-1.15 (s390x) xorg-x11-libXau-7.4-1.15 (s390x) xorg-x11-libXext-32bit-7.4-1.16.21 (s390x) xorg-x11-libXext-7.4-1.16.21 (s390x) xorg-x11-libs-32bit-7.4-8.26.32.1 (s390x) xorg-x11-libs-7.4-8.26.32.1 (s390x) xorg-x11-libxcb-32bit-7.4-1.20.34 (s390x) xorg-x11-libxcb-7.4-1.20.34 (s390x) OCFS2 1.4 (For Oracle RAC only) </pre>

4.10 Additional Drivers and Software Packages for Linux

You are not required to install additional drivers and packages, but you may choose to install or configure drivers and packages in the following list:

- [Installation Requirements for Open Database Connectivity](#)
- [Installation Requirements for PAM on Linux](#)
- [Installation Requirements for OCFS2](#)
- [Installation Requirements for Oracle Messaging Gateway](#)
- [Installation Requirements for Lightweight Directory Access Protocol](#)
- [Installation Requirements for Programming Environments for Linux](#)
- [Installation Requirements for Web Browsers](#)

4.10.1 Installation Requirements for Open Database Connectivity

Review the following sections if you plan to install Open Database Connectivity (ODBC):

- [About ODBC Drivers and Oracle Database](#)
- [Installing ODBC Drivers for Linux x86-64](#)

4.10.1.1 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.

4.10.1.2 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

4.10.2 Installation Requirements for PAM on Linux

Review the following sections to install PAM:

- [About PAM and Login Authentication](#)
- [Installing PAM Library](#)

4.10.2.1 About PAM and Login Authentication

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.

4.10.2.2 Installing PAM Library

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

4.10.3 Installation Requirements for OCFS2

Review the following sections to install OCFS2

- [About OCFS2 and Shared Storage](#)
- [Installing OCFS2](#)

4.10.3.1 About OCFS2 and Shared Storage

Oracle Cluster File System 2 (OCFS2) is a POSIX-compliant general purpose shared disk cluster file system for Linux. You can use OCFS2 with Oracle Grid Infrastructure. However, you are not required to use OCFS2. OCFS2 is supported for this release only with Oracle Linux 5 and Oracle Linux 6.

On Linux, OCFS2 is supported for use with Regular Cluster deployments for OCR and voting files. OCFS2 is not supported for Grid homes, and is not supported for Oracle Flex Cluster deployments.

4.10.3.2 Installing OCFS2

OCFS2 Release 2.1.6 is included with the Unbreakable Enterprise Kernel for Oracle Linux available with Oracle Linux 5 and Oracle Linux 6.

See the OCFS2 project page for additional information:

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

4.10.4 Installation Requirements for Oracle Messaging Gateway

Review the following sections to install Oracle Messaging Gateway

- [About Oracle Messaging Gateway](#)
- [Installing Oracle Messaging Gateway](#)

4.10.4.1 About Oracle Messaging Gateway

Oracle Messaging Gateway is a feature of Oracle Database. It enables communication between applications based on non-Oracle messaging systems and Oracle Streams Advanced Queuing.

Oracle Messaging Gateway supports the integration of Oracle Streams Advanced Queuing (AQ) with applications based on WebSphere and TIBCO Rendezvous. For information on supported versions, see *Oracle Database Advanced Queuing User's Guide*

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

4.10.4.2 Installing Oracle Messaging Gateway

Oracle Messaging Gateway is installed with the Enterprise Edition of Oracle Database.

If you require a CSD for IBM WebSphere MQ, then see the following website for download and installation information:

<http://www.ibm.com>

4.10.5 Installation Requirements for Lightweight Directory Access Protocol

Review the following sections to install Lightweight Directory Access Protocol:

- [About LDAP and Oracle Plug-ins](#)
- [Installing the LDAP Package](#)

4.10.5.1 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 to use features requiring LDAP, including the Oracle Database scripts `odisrvreg` and `oidca` for Oracle Internet Directory, or `schemasync` for third-party LDAP directories.

4.10.5.2 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.

4.10.6 Installation Requirements for Programming Environments for Linux

Review the following sections to install programming environments:

- [About Programming Environments and Oracle Database](#)
- [Configuring Support for Programming Environments](#)

4.10.6.1 About Programming Environments and Oracle Database

Oracle Database supports multiple programming languages for application development in different environments. Some languages require that you install additional compiler packages for the operating system.

Programming environments are options. They are not required for Oracle Database.

See Also : *Oracle Database Advanced Application Developer's Guide* for an overview of programming environments

4.10.6.2 Configuring Support for Programming Environments

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

- [Requirements for Programming Environments for x86-64 Linux](#)
- [Requirements for Programming Environments for IBM: Linux on System z](#)
- [Requirements for Programming Environments for Linux on SPARC](#)

Table 4–10 *Requirements for Programming Environments for x86-64 Linux*

Programming Environments	Support Requirements
Java Database Connectivity, Oracle Call Interface (OCI)	JDK 6 (Java SE Development Kit release 1.6.0_37 or later updates of 1.6) with the JNDI extension with Oracle Java Database Connectivity. JDK 1.6 is installed with this release.
Oracle C++ Oracle C++ Call Interface Pro*C/C++ Oracle XML Developer's Kit (XDK)	Intel C/C++ Compiler 12.05 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	Micro Focus Server Express 5.1

Table 4–11 Requirements for Programming Environments for IBM: Linux on System z

Programming Environments	Support Requirements
Java Database Connectivity/Oracle Call Interface (OCI)	<ul style="list-style-type: none"> ■ JDK 6 (1.6.0 SR12) ■ JDK 7 (1.7.0) JDK 1.6 is installed with this release.
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	Micro Focus Server Express 5.1

Table 4–12 Requirements for Programming Environments for Linux on SPARC

Programming Environments	Support Requirements
Java Database Connectivity/Oracle Call Interface (OCI)	JDK 8 (1.8)
Oracle C++ Oracle C++ Call Interface Pro*C/C++ Oracle XML Developer's Kit (XDK)	Oracle Solaris Studio 12.5 for Linux/SPARC (formerly Sun Studio): Studio 12.5 Sun C 5.14 Linux_sparc 2016/02/19Sun

4.10.7 Installation Requirements for Web Browsers

Web browsers are required to use Oracle Enterprise Manager Database Express and Oracle Enterprise Manager Cloud Control. Web browsers must support Java Script, and the HTML 4.0 and CSS 1.0 standards. For a list of browsers that meet these requirements, see the Oracle Enterprise Manager certification matrix on My Oracle Support:

<https://support.oracle.com>

See Also: *Oracle Enterprise Manager Cloud Control Basic Installation Guide* for information on accessing the Oracle Enterprise Manager certification matrix

4.11 Checking the Software Requirements

To ensure that the system meets these requirements, follow these steps:

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

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

2. To determine which distribution and version of IBM: Linux on System z is installed, enter the one of the following commands:

```
# cat /etc/SuSE-release
# cat /etc/redhat-release
```

```
# lsb_release -id
```

3. To determine whether the required kernel errata is installed, enter the following command:

```
# uname -r
```

The following is sample output displayed by running this command 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.

4. 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++ elfutils 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.

4.12 Installing the cvuqdisk RPM for Linux

If you do not use an Oracle Preinstallation RPM, 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`).

To install the `cvuqdisk` RPM, complete the following procedure:

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, 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 will own `cvuqdisk`, typically `oinstall`. For example:

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

6. In the directory where you have saved the `cvuqdisk rpm`, use the following command to install the `cvuqdisk` package:

```
# rpm -iv package
```

For example:

```
# rpm -iv cvuqdisk-1.0.9-1.rpm
```

4.13 Checking Shared Memory File System Mount on Linux

Ensure that the `/dev/shm` mount area is of type `tmpfs` and is mounted with the following options:

- With `rw` and `exec` permissions set on it
- Without `noexec` or `nosuid` set on it

Use the following procedure to check the shared memory file system:

1. Check current mount settings. For example:

```
$ more /etc/fstab |grep "tmpfs"
tmpfs          /dev/shm      /tmpfs        defaults      0 0
```

2. If necessary, change mount settings. For example, log in as `root`, open the `/etc/fstab` file with a text editor, and modify the `tmpfs` line:

```
tmpfs          /dev/shm      /tmpfs        rw,exec      0 0
```

See Also: *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for more information about shared memory mounts

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
```

In the preceding example, `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 one of 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
```

4.15 Setting the Disk I/O Scheduler on Linux

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

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

```
# echo deadline > /sys/block/${ASM_DISK}/queue/scheduler
```

4.16 Setting Network Time Protocol for Cluster 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 environment variable setting of the Grid installation owner on the node where 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)
- 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.

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.

To deactivate the NTP service, you must stop the existing ntpd service, disable it from the initialization sequences and remove the ntp.conf file. To complete these step on Oracle Linux, and Asianux systems, run the following commands as the root user

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

Also 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.

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

```
$ crsctl check ctss
```

If you are using NTP, and you prefer to continue using it instead of Cluster Time Synchronization Service, then you need to modify the NTP configuration to set the `-x` flag, which prevents time from being adjusted backward. Restart the network time protocol daemon after you complete this task.

To do this, on Oracle Linux, Red Hat Linux, and Asianux systems, edit the `/etc/sysconfig/ntpd` file to add the `-x` flag, as in the following example:

```
# Drop root to id 'ntp:ntp' by default.
OPTIONS="-x -u ntp:ntp -p /var/run/ntpd.pid"

# Set to 'yes' to sync hw clock after successful ntpdate
SYNC_HWCLOCK=no

# Additional options for ntpdate
NTPDATE_OPTIONS=""
```

Then, restart the NTP service:

```
# /sbin/service ntpd restart
```

On SUSE systems, modify the configuration file `/etc/sysconfig/ntp` with the following settings:

```
NTPD_OPTIONS="-x -u ntp"
```

Restart the daemon using the following command:

```
# service ntpd restart
```

4.17 Using Automatic SSH Configuration During Installation

To install Oracle software, Secure Shell (SSH) connectivity should be set up between all cluster member nodes. 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 the Oracle Universal Installer (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
```

See Also: [Section 6.2.5, "Preventing Installation Errors Caused by Terminal Output Commands"](#) for information about how to remove `stty` commands in user profiles

Configuring Networks for Oracle Grid Infrastructure and Oracle RAC

Review the following sections to check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Grid Infrastructure for a cluster installation.

This chapter contains the following topics:

- [Network Interface Hardware Requirements](#)
- [IP Interface Configuration Requirements](#)
- [Private Interconnect Redundant Network Requirements](#)
- [IPv4 and IPv6 Protocol Requirements](#)
- [Oracle Grid Infrastructure IP Name and Address Requirements](#)
- [About Oracle Flex ASM Clusters Networks](#)
- [Broadcast Requirements for Networks Used by Oracle Grid Infrastructure](#)
- [Multicast Requirements for Networks Used by Oracle Grid Infrastructure](#)
- [Domain Delegation to Grid Naming Service](#)
- [Configuration Requirements for Oracle Flex Clusters](#)
- [Grid Naming Service Standard Cluster Configuration Example](#)
- [Manual IP Address Configuration Example](#)
- [Network Interface Configuration Options](#)
- [Multiple Private Interconnects and Oracle Linux](#)

See Also: 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>

5.1 Network Interface Hardware Requirements

The following is a list of requirements for network configuration:

- Each node must have at least two network adapters or network interface cards (NICs): one for the public network interface, and one for the private network interface (the interconnect).

When you upgrade a node to Oracle Grid Infrastructure 11g Release 2 (11.2.0.2) and later, the upgraded system uses your existing network classifications.

To configure multiple public interfaces, use a third-party technology for your platform to aggregate the multiple public interfaces before you start installation, and then select the single interface name for the combined interfaces as the public interface. Oracle recommends that you do not identify multiple public interface names during Oracle Grid Infrastructure installation. Note that if you configure two network interfaces as public network interfaces in the cluster without using an aggregation technology, the failure of one public interface on a node does not result in automatic VIP failover to the other public interface.

Oracle recommends that you use the Redundant Interconnect Usage feature to make use of multiple interfaces for the private network. However, you can also use third-party technologies to provide redundancy for the private network.

Note: Redundant Interconnect Usage requires a complete Oracle Grid Infrastructure and Oracle Database Release 2 (11.2.0.2) or higher stack. Earlier release Oracle Databases cannot use this feature, and must use third-party NIC bonding technologies. If you consolidate different database releases in one cluster, and use databases before Oracle Database 11g Release 2 (11.2.0.2), then you may require both technologies.

- For the public network, each network adapter must support TCP/IP.
- For the private network, the interface must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (minimum requirement 1 Gigabit Ethernet).

Note: UDP is the default interface protocol for Oracle RAC and Oracle Clusterware. You must use a switch for the interconnect. Oracle recommends that you use a dedicated switch.

Oracle does not support token-rings or crossover cables for the interconnect.

5.2 IP Interface Configuration Requirements

- For clusters using single interfaces for private networks, each node's private interface for interconnects must be on the same subnet, and that 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].

For clusters using Redundant Interconnect Usage, each private interface should be on a different subnet. However, 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.

- For the private network, the 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`.

5.3 Private Interconnect Redundant Network Requirements

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. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2). If you use the Oracle Clusterware Redundant Interconnect feature, then you must use IPv4 addresses for the interfaces.

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.

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

5.4 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.

Cluster member node interfaces can be configured to use IPv4, IPv6, or both types of Internet protocol addresses. However, be aware of the following:

- **Configuring public VIPs:** 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. For example, 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 must configure the private network as an IPv4 network. IPv6 addresses are not supported for the interconnect.
- **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.

See Also:

- <http://www.ietf.org/rfc/rfc2732.txt> for RFC 2732, and information about IPv6 notational representation
- <http://www.ietf.org/rfc/rfc3513.txt> for RFC 3513, and information about proper IPv6 addressing
- <http://www.ietf.org/rfc/rfc2462.txt> for RFC 2462, and information about IPv6 Stateless Address Autoconfiguration protocol
- *Oracle Database Net Services Administrator's Guide* for more information about network communication and IP address protocol options

5.5 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. The Oracle Grid Naming Service is used with large clusters to ease network administration costs.

This section contains the following topics:

- [About Oracle Grid Infrastructure Name Resolution Options](#)
- [Cluster Name and SCAN Requirements](#)
- [IP Name and Address Requirements For Grid Naming Service \(GNS\)](#)

- [IP Name and Address Requirements for Standard Cluster Manual Configuration](#)

5.5.1 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.

You can 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.

Note: Oracle recommends that you use a static host name for all non-VIP server node public host names.

Public IP addresses and virtual IP addresses must be in the same subnet.

5.5.2 Cluster Name and SCAN 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.

If you configure a Standard cluster, and choose a Typical install, then the SCAN is also the name of the cluster. In that case, the SCAN must meet the requirements for a cluster name. The SCAN can be no longer than 15 characters.

In an Advanced installation, 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 name carefully. After installation, you can only change the cluster name by reinstalling Oracle Grid Infrastructure.

5.5.3 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:

- For Standard Clusters: If you have vendor clusterware installed, then you cannot choose to use GNS, because the vendor clusterware does not support it. Vendor clusterware is not supported with Oracle Flex Cluster configurations.
 - You cannot use GNS with another multicast DNS. To use GNS, disable any third party mDNS daemons on your system.
-
-

5.5.4 IP Name and Address Requirements For Multi-Cluster GNS

Review the following requirements for using Multi-cluster GNS:

- [About Multi-Cluster GNS Networks](#)
- [Configuring GNS Server Clusters](#)
- [Configuring GNS Client Clusters](#)
- [Creating and Using a GNS Client Data File](#)

5.5.4.1 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.

To provide networking service, Multi-cluster 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.

5.5.4.2 Configuring GNS Server Clusters

To use this option, 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.

5.5.4.3 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.

5.5.4.4 Creating and Using a GNS Client Data File

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
```

For example:

```
$ srvctl export gns -clientdata /home/grid/gns_client_data
```

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.

After you have completed the GNS client cluster installation, you must run the following command on one of the GNS server cluster members to start GNS service, where *path_to_file* is the name and path location of the GNS client data file:

```
srvctl add gns -clientdata path_to_file
```

For example:

```
$ srvctl add gns -clientdata /home/grid/gns_client_data
```

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

5.5.5 IP Name and Address Requirements for Standard Cluster Manual Configuration

If you do not enable GNS, then you must configure static cluster node names and addresses before starting installation.

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.

The cluster must have the following names and addresses:

- 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
- 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
- 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 ("_")
- 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: In a Typical installation, the SCAN you provide is also the name of the cluster, so the SCAN name must meet the requirements for a cluster name. In an Advanced installation, 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.

See Also: [Appendix E, "Understanding Network Addresses"](#) for more information about network addresses

5.5.6 Confirming the DNS Configuration for SCAN

You can use the `nslookup` command to confirm that the DNS is correctly associating the SCAN with the addresses. For example:

```
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.

Note: 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 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. Each Oracle Flex ASM cluster has its own name that is globally unique within the enterprise.

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.

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.

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for more information about Oracle Flex Clusters
- *Oracle Automatic Storage Management Administrator's Guide* for more information about Oracle Flex ASM

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 Flex ASM cluster network has the following requirements and characteristics:

- The ASM network can be configured during installation, or configured or modified after installation.

Cluster nodes can be configured as follows:

- Oracle Flex ASM cluster Hub Nodes, 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 ASM disks.
 - Run a Direct ASM client process.
 - Run an ASM Filter Driver, part of whose function is to provide cluster fencing security for the Oracle Flex ASM cluster.
 - Access the 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 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.
- Oracle Flex ASM cluster Leaf Nodes, with the following characteristics:
 - Use Indirect access to the ASM disks, where I/O is handled as a service for the client on a Hub Node.
 - Submit disk service requests through the ASM network.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about Oracle Flex ASM clusters

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 ARP. Highly Available IP (HAIP) addresses do not require ARP on the public network, but for VIP failover, you will 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. 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.

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. 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`.

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for more information about GNS
- [Section 5.5.2, "Cluster Name and SCAN Requirements"](#) for information about choosing network identification names

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).

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about GNS options, delegation, and public networks

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/resolve.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
```

Note: Be aware that use of NIS is a frequent source of problems when doing cable pull tests, as host name and username resolution can fail.

5.10 Configuration Requirements for Oracle Flex Clusters

Review the following information if you intend to configure an Oracle Flex Cluster:

- [General Requirements for Oracle Flex Cluster Configuration](#)
- [Oracle Flex Cluster DHCP-Assigned Virtual IP \(VIP\) Addresses](#)
- [Oracle Flex Cluster Manually-Assigned Addresses](#)

5.10.1 General Requirements for Oracle Flex Cluster Configuration

Note the following 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.

5.10.2 Oracle Flex Cluster DHCP-Assigned Virtual IP (VIP) Addresses

If you choose to configure DHCP-assigned VIPs, then during installation select one of the following options to configure cluster node VIP names for both Hub and Leaf Nodes:

- **Manual Names:** Enter the node name and node VIP name for each cluster member node (for example, linnode1; linnode1-vip; linnode2; linnode2-vip; and so on) to be assigned to the VIP addresses delegated to cluster member nodes through DHCP, and resolved by DNS. Manual names must confirm with the RFC 952 standard, which allows alphanumeric characters and hyphens ("-"), but does not allow underscores ("_").
- **Automatically Assigned Names:** Select **Auto Assigned** to allow the installer to assign names to VIP addresses generated through DHCP automatically. using the pattern *name#* and *name#-vip*, where *name* is the cluster name and # is an automatically assigned number. Addresses are assigned through DHCP, and resolved by GNS.

5.10.3 Oracle Flex Cluster 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.

You can create manual addresses using alphanumeric strings. For example, the following strings are examples of acceptable names: mycloud001nd; mycloud046nd; mycloud046-vip; mycloud348nd; mycloud784-vip.

5.11 Grid Naming Service Standard Cluster Configuration Example

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 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 Example Network

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 ¹	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 ¹	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.cluster01.example.com	virtual	192.0.2.201	DHCP	GNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster-scan.cluster01.example.com	virtual	192.0.2.202	DHCP	GNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster-scan.cluster01.example.com	virtual	192.0.2.203	DHCP	GNS

¹ Node host names may resolve to multiple addresses, including VIP addresses currently running on that host.

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. Also, 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 ¹	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 ¹	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

¹ Node host names may resolve to multiple addresses.

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

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

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
```

Oracle Linux 5.6 (Oracle Linux 5 Update 6) includes a fix using `initscripts-8.45.33-1.0.4.el5.i386.rpm`, which sets the kernel parameter `net.ipv4.conf.default.rp_filter` to 2 (relaxed mode). For that reason, after you apply the Unbreakable Linux kernel on top of Oracle Linux 5.6, you may not need to make manual changes, because the `rp_filter` value of all NICs is set to 2. If you require more strict reverse path filtering on the public network, then set the public NIC `rp_filter` to 1.

See Also: 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/CSP/main/article?cmd=show&type=NOT&id=1286796.1>

Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle RAC

This chapter describes the users, groups user environment and management environment settings to complete before you install Oracle Grid Infrastructure for a Cluster and Oracle Real Application Clusters.

This chapter contains the following topics:

- [Creating Groups, Users and Paths for Oracle Grid Infrastructure](#)
- [Configuring Grid Infrastructure Software Owner User Environments](#)
- [Enabling Intelligent Platform Management Interface \(IPMI\)](#)
- [Determining Root Script Execution Plan](#)

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 a software owner for Oracle Grid Infrastructure.

- [Determining If the Oracle Inventory and Oracle Inventory Group Exists](#)
- [Creating the Oracle Inventory Group If an Oracle Inventory Does Not Exist](#)
- [Creating the Oracle Grid Infrastructure User](#)
- [About the Oracle Base Directory for the Grid User](#)
- [About the Oracle Home Directory for Oracle Grid Infrastructure Software](#)
- [Creating the Oracle Home and Oracle Base Directory](#)
- [About Job Role Separation Operating System Privileges Groups and Users](#)
- [Example of Creating Minimal Groups, Users, and Paths](#)
- [Example of Creating Role-allocated Groups, Users, and Paths](#)

Note: During an Oracle Grid Infrastructure installation, both Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) are installed. You no longer can have separate Oracle Clusterware installation owners and Oracle ASM installation owners.

6.1.1 Determining If the Oracle Inventory and Oracle Inventory Group Exists

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. An `oraInst.loc` file has contents similar to the following:

```
inventory_loc=central_inventory_location
inst_group=group
```

In the preceding example, `central_inventory_location` is the location of the Oracle central inventory, and `group` is the name of the group that has permissions to write to the central inventory (the OINSTALL system privilege).

For Oracle Grid Infrastructure installations, the central inventory must be on local storage on the node.

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.

To determine if you have an Oracle central inventory directory (`oraInventory`) on your system:

Enter the following command:

```
# more /etc/oraInst.loc
```

If the `oraInst.loc` file exists, then the output from this command is similar to the following:

```
inventory_loc=/u01/app/oracle/oraInventory
inst_group=oinstall
```

In the previous output example:

- The `inventory_loc` group shows the location of the Oracle Inventory.
- The `inst_group` parameter shows the name of the Oracle Inventory group (in this example, `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
```

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 by entering a command similar to the following:

```
# /usr/sbin/groupadd -g 54321 oinstall
```

The preceding command creates the `oraInventory` group `oinstall`, with the group ID number 54321. Members of the `oraInventory` group are granted privileges to write to the Oracle central inventory (`oraInventory`), and other system privileges for Oracle installation owner users.

By default, if an `oraInventory` group does not exist, then the installer lists the primary group of the installation owner for the Oracle Grid Infrastructure for a cluster software as the `oraInventory` group. Ensure that this group is available as a primary group for all planned Oracle software installation owners.

Note: Group and user IDs must be identical on all nodes in the cluster. Check to make sure 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 Creating the Oracle Grid Infrastructure User

You must create a software owner for Oracle Grid Infrastructure 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. A user created to own either all Oracle installations, or only Oracle database installations, is called the `oracle` user.

6.1.3.1 Understanding 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/12.1.0/grid` is the Oracle Grid Infrastructure home:

```
$ ORACLE_HOME=/u01/app/12.1.0/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.

6.1.3.2 Determining if an Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` or `grid` exists, enter a command similar to the following (in this case, to determine if `oracle` exists):

```
# id oracle
```

If the user exists, then the output from this command is similar to the following:

```
uid=54321(oracle) gid=54321(oinstall) groups=54322(dba),54323(oper)
```

Determine whether you want to use the existing user, or create another user. The user and group ID numbers must be the same on each node you intend to make a cluster member node.

If you are using the Oracle Preinstallation RPM to provision your Linux operating system for an Oracle Grid Infrastructure and Oracle Database installation, then it has configured for you the Oracle database installation owner (`oracle`), an Oracle Inventory group (`oinstall`), and an Oracle administrative privileges group (`dba`).

To use an existing user as an installation owner for this installation, ensure that the user's primary group is the Oracle Inventory group (`oinstall`).

6.1.3.3 Creating or Modifying an Oracle Software Owner User for Oracle Grid Infrastructure

If the Oracle software owner (`oracle`, `grid`) user does not exist, or if you require a new Oracle software owner user, then create it. If you want to use an existing user account, then modify it to ensure that the user IDs (UID) and group IDs (GID) are the same on each cluster member node.

Oracle recommends that you do not use the defaults on each node, because UIDs and GIDs assigned by default are likely to be different on each node. Instead, determine specifically named and numbered UIDs and GIDs, and confirm that they are unused on any node before you create or modify groups and users. Oracle strongly recommends that you confirm that you have identical user configuration on each node you intend to make a cluster member node before you start installation, and that the

UID and GIDs do not need to be changed. If you need to change UIDs and GIDs for Oracle software users and groups, then you must reinstall the software.

Oracle does not support changing the UID or GID or group memberships for a user account that you have previously used to install and configure Oracle RAC or Oracle Grid Infrastructure. Oracle does not support changing the ownership of an existing Oracle Database home from one Oracle user to a different user.

If you must modify an existing Grid installation owner after previously installing Oracle Grid Infrastructure (for example, if you want to modify an existing Oracle Database user account before you install Oracle RAC, to distinguish between the Oracle Grid Infrastructure owner and the Oracle RAC Oracle Database owner user accounts), then you must stop and start Oracle Clusterware on each node (in a rolling fashion) to pick up the changes made to the user account that owns Oracle Grid Infrastructure.

The following procedures use `grid` as the name of the Oracle software owner, and `asmadmin` as the OSASM group. To create separate system privileges groups to separate administration privileges, complete group creation before you create the user, as described in [Section 6.1.7, "About Job Role Separation Operating System Privileges Groups and Users,"](#) on page 6-8.

1. To create a grid installation owner account where you have an existing system privileges group (in this example, `dba`), whose members you want to have granted the SYSASM privilege to administer the Oracle ASM instance, enter a command similar to the following:

```
# /usr/sbin/useradd -u 54322 -g oinstall -G dba grid
```

In the preceding command:

- The `-u` option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, Oracle recommends that you specify a number. You must make note of the user ID number of the user you create for Oracle Grid Infrastructure, as you require it later during preinstallation. You must use the same user ID number for this user on all nodes of the cluster.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group. For example: `oinstall`.
- The `-G` option specified the secondary group, which in this example is `dba`.

The secondary groups must include the OSASM group, whose members are granted the SYSASM privilege to administer the Oracle ASM instance. You can designate a unique group for the SYSASM system privilege, separate from database administrator groups, or you can designate one group as the OSASM and OSDBA group, so that members of that group are granted the SYSASM and SYSDBA privilege to grant system privileges to administer both the Oracle ASM instances and Oracle Database instances. In code examples, this group is `asmadmin`.

If you are creating this user to own both Oracle Grid Infrastructure and an Oracle Database installation, then this user must have the OSDBA for ASM group as a secondary group. In code examples, this group name is `asmdba`. Members of the OSDBA for ASM group are granted access to Oracle ASM storage. You must create an OSDBA for ASM group if you plan to have multiple databases accessing Oracle ASM storage, or you must use the same group as the OSDBA for all databases, and for the OSDBA for ASM group.

You are also prompted during installation to assign operating system groups for several other Oracle Database system administrative privileges.

Use the `usermod` command to change existing user ID numbers and groups.

For example:

```
# id oracle
uid=1101(oracle) gid=1000(oinstall) groups=1200(dba)
# /usr/sbin/usermod -u 54421 -g 54331 -G 54321,54322,54327 oracle
# id oracle
uid=54321(oracle) gid=54321(oinstall) groups=54321(oinstall),54322(dba),54327
(asmdba)
```

2. Set the password of the user that will own Oracle Grid Infrastructure. For example:

```
# passwd grid
```

3. Repeat this procedure on all of the other nodes in the cluster.

Note: If necessary, contact your system administrator before using or modifying an existing user.

6.1.4 About the Oracle Base Directory for the Grid User

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
```

6.1.5 About the Oracle Home Directory for Oracle Grid Infrastructure Software

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/12.1.0/grid
```

During installation, ownership of the entire path to the Grid home is changed to root. (/u01, /u01/app, /u01/app/12.1.0, /u01/app/12.1.0/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:

```
/u01/12.1.0/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.

See Also: *Oracle Database Installation Guide* for details about Optimal Flexible Architecture (OFA) guidelines

6.1.6 Creating the Oracle Home and Oracle Base Directory

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/12.1.0/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.

See Also: [Appendix E, "Oracle Grid Infrastructure for a Cluster Installation Concepts"](#) for more details about Oracle base and Oracle home directories

6.1.7 About Job Role Separation Operating System Privileges 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.

With Oracle Database job role separation, each Oracle Database installation has separate operating system groups to provide authentication for system privileges on that Oracle Database, so multiple databases can be installed on the cluster without sharing operating system authentication for system privileges. In addition, each Oracle software installation is owned by a separate installation owner, to provide operating system user authentication 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 Clusterware resources.

With Oracle Grid Infrastructure job role separation, Oracle ASM has separate operating system groups that provides operating system authentication for Oracle ASM system privileges for storage tier administration. This operating system authentication is separated from Oracle Database operating system authentication. In addition, the Oracle Grid Infrastructure installation owner provides operating system user authentication for modifications to Oracle Grid Infrastructure binaries.

During the Oracle Database installation, Oracle Universal Installer (OUI) prompts you to specify the name of the OSDBA, OSOPER, OSBACKUPDBA, OSDGDBA and OSKMDBA groups. Members of these groups are granted operating system authentication for the set of database system privileges each group authorizes. Oracle recommends that you create different operating system groups for each set of system privileges.

Note: This configuration is optional, to restrict user access to Oracle software by responsibility areas for different administrator users.

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).

To simplify using the defaults for Oracle tools such as Cluster Verification Utility, if you do choose to use a single operating system group to grant all system privileges and the right to write to the oraInventory, then that group name should be oinstall.

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.

See Also:

- *Oracle Database Administrator's Guide* for more information about planning for system privileges authentication
- *Oracle Automatic Storage Management Administrator's Guide* for more information about Oracle ASM operating system authentication

6.1.8 Descriptions of Job Role Separation Groups and Users

This section contains the following topics:

- [Oracle Software Owner For Each Oracle Software Product](#)
- [Standard Oracle Database Groups for Job Role Separation](#)
- [Extended Oracle Database Groups for Job Role Separation](#)
- [Oracle ASM Groups for Job Role Separation](#)

6.1.8.1 Oracle Software Owner For Each Oracle Software Product

Oracle recommends that you create one software owner to own each Oracle software product (typically, `oracle`, for the database software owner user, and `grid` for Oracle Grid Infrastructure).

You must create at least one software owner the first time you install Oracle software on the system. This user owns the Oracle binaries of the Oracle Grid Infrastructure software, and you can also make this user the owner of the Oracle Database or Oracle RAC binaries.

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, and OSKMDBA groups as secondary groups. In Oracle documentation, when Oracle software owner users are referred to, they are called `oracle` users.

Oracle recommends that you create separate software owner users to own each Oracle software installation. Oracle particularly recommends that you do this if you intend to install multiple databases on the system.

In Oracle documentation, a user created to own the Oracle Grid Infrastructure binaries is called the `grid` user. This user owns both the Oracle Clusterware and Oracle Automatic Storage Management binaries.

See Also: *Oracle Automatic Storage Management Administrator's Guide* and *Oracle Database Administrator's Guide* for more information about operating system groups and system privileges authentication

6.1.8.2 Standard Oracle Database Groups for Job Role Separation

The following is a list of standard Oracle Database groups. These groups provide operating system authentication for database administration system privileges:

- **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.

To specify a group name other than the default `dba` group, you must either choose the Advanced installation type to install the software, or you start Oracle Universal Installer (OUI) as a user that is not a member of this group. If you start OUI with a user that is not a member of a group called `dba`, then OUI prompts you to specify the name of the `OSDBA` group.

Members of the `OSDBA` group formerly were granted `SYSASM` privilege on Oracle ASM instances, including mounting and dismounting disk groups. This privileges grant is removed with Oracle Grid Infrastructure 11g Release 2 (11.2), if different operating system groups are designated as the `OSDBA` and `OSASM` groups. If the same group is used for both `OSDBA` and `OSASM`, then the privileges are retained.

- **OSOPER group for Oracle Database** (typically, `oper`)

You can choose to create this group if you want a separate group of operating system users to have a limited set of database administrative privileges for starting up and shutting down the database (the `SYSOPER` privilege).

6.1.8.3 Extended Oracle Database Groups for Job Role Separation

Starting with Oracle Database 12c Release 1 (12.1), in addition to the `OSOPER` privileges to start up and shut down the database, you can create new administrative privileges that are more task-specific and less privileged than the `OSDBA/SYSDBA` system privileges to support specific administrative privileges tasks required for everyday database operation. Users granted these system privileges are also authenticated through operating system group membership.

You do not have to create these specific group names, but during installation you are prompted to provide operating system groups whose members are granted access to these system privileges. You can assign the same group to provide authentication for

these privileges, but Oracle recommends that you provide a unique group to designate each privilege.

The OSDBA subset 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 `SYSDBG` privilege).
- **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).

6.1.8.4 Oracle ASM Groups for Job Role Separation

The `SYSASM`, `SYSOPER` for ASM, and `SYSDBA` for ASM system privileges enables the separation of the Oracle ASM storage administration privileges from `SYSDBA`. Members of operating systems you designate are granted the system privileges for these roles. Select separate operating system groups as the operating system authentication groups for privileges on Oracle ASM.

Before you start OUI, create the following OS groups and users for Oracle ASM, whose members are granted the corresponding SYS privileges:

- **OSASM Group for Oracle ASM Administration** (typically `asmadmin`)
Create this group as a separate group if you want to have 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.
- **OSDBA for ASM Database Administrator group for ASM**, typically `asmdba`)
Members of the ASM Database Administrator group (OSDBA for ASM) are granted read and write access to files managed by Oracle ASM. The Oracle Grid Infrastructure installation owner and all Oracle Database software owners must be a member of this group, and all users with OSDBA membership on databases that

have access to the files managed by Oracle ASM must be members of the OSDBA group for ASM.

- **OSOPER for ASM Group for ASM Operators** (OSOPER for 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 ASM 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.

To use the Oracle ASM Operator group to create an Oracle ASM administrator group with fewer privileges than the default `asmadmin` group, then you must choose the Advanced installation type to install the software. In this case, OUI prompts you to specify the name of this group. In code examples, this group is `asmoper`.

6.1.9 Creating Job Role Separation Operating System Privileges Groups and User

The following sections describe how to create the required operating system user and group for Oracle Grid Infrastructure and Oracle Database:

- [Creating the OSDBA Group to Prepare for Database Installations](#)
- [Creating an OSOPER Group for Database Installations](#)
- [Creating the OSASM Group](#)
- [Creating the OSOPER for ASM Group](#)
- [Creating the OSDBA for ASM Group for Database Access to Oracle ASM](#)
- [Oracle Software Owner User Installation Tasks](#)
- [Creating Identical Database Users and Groups on Other Cluster Nodes](#)

6.1.9.1 Creating the OSDBA Group to Prepare for Database Installations

If you intend to install Oracle Database to use with the Oracle Grid Infrastructure installation, then 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

If the OSDBA group does not exist, or if you require a new OSDBA group, then create it. Use the group name `dba` unless a group with that name already exists. For example:

```
# /usr/sbin/groupadd -g 54322 dba
```

6.1.9.2 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. To use an OSOPER group, then you must 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
```

6.1.9.3 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
```

6.1.9.4 Creating the OSOPER for ASM Group

Create an OSOPER for ASM group if you want to identify a group of operating system users, such as database administrators, whom you want to grant a limited set of Oracle ASM storage tier administrative privileges, including the ability to start up and shut down the Oracle ASM storage. For most installations, it is sufficient to create only the OSASM group, and provide that group as the OSOPER for ASM group during the installation interview.

If the OSOPER for ASM group does not exist, or if you require a new OSOPER for ASM group, then create it. Use the group name `asmoper` unless a group with that name already exists. For example:

```
# groupadd -g 54328 asmoper
```

6.1.9.5 Creating the OSDBA for ASM Group for Database Access to Oracle ASM

You must create an OSDBA for ASM group to provide access to the Oracle ASM instance. This is necessary if OSASM and OSDBA are different groups.

If the OSDBA for ASM group does not exist, or if you require a new OSDBA for ASM group, then create it. Use the group name `asmdba` unless a group with that name already exists. For example:

```
# groupadd -g 54327 asmdba
```

6.1.9.6 Oracle Software Owner User Installation Tasks

This section contains information about the Oracle software owner user, which is typically the user that owns Oracle Database or other Oracle application software. This section contains the following topics:

- [About the Oracle Software Owner User](#)
- [Determining if an Oracle Software Owner User Exists](#)
- [Creating an Oracle Software Owner User](#)
- [Modifying an Existing Oracle Software Owner User](#)

6.1.9.6.1 About the Oracle Software Owner User You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user exists, but you want to use a different operating system user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation.
- If you have created an Oracle software owner for Oracle Grid Infrastructure, such as `grid`, and you want to create a separate Oracle software owner for Oracle Database software, such as `oracle`.

6.1.9.6.2 Determining if an Oracle Software Owner User Exists To determine whether an Oracle software owner user named `oracle` or `grid` exists, enter a command similar to the following (in this case, to determine if `oracle` exists):

```
# id oracle
```

If the user exists, then the output from this command is similar to the following:

```
uid=54321(oracle) gid=54321(oinstall) groups=54322(dba),54323(oper)
```

Determine whether you want to use the existing user, or create another user. To use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

- To modify an existing user, see [Section 6.1.9.6.4, "Modifying an Existing Oracle Software Owner User."](#)
- To create a user, see [Section 6.1.9.6.3, "Creating an Oracle Software Owner User."](#)

Note: If necessary, contact your system administrator before using or modifying an existing user.

Oracle recommends that you do not use the UID and GID defaults on each node, as group and user IDs likely will be different on each node. Instead, provide common assigned group and user IDs, and confirm that they are unused on any node before you create or modify groups and users.

6.1.9.6.3 Creating an Oracle Software Owner User If the Oracle software owner user does not exist, or if you require a new Oracle software owner user, then create it. Use the user name `oracle` unless a user with that name already exists.

To create an `oracle` user:

1. Enter a command similar to the following:

```
# useradd -u 54322 -g oinstall -G dba,asmdba oracle
```

In the preceding command:

- The `-u` option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, you must make note of the `oracle` user ID number, as you require it later during preinstallation.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group. For example: `oinstall`.
- The `-G` option specifies the secondary groups, which must include the OSDBA group, the OSDBA for ASM group, and, if required, the OSOPER for ASM group. For example: `dba, asmdba, or dba, asmdba, asmoper`.

2. Set the password of the oracle user:

```
# passwd oracle
```

6.1.9.6.4 Modifying an Existing Oracle Software Owner User If the `oracle` user exists, but its primary group is not `oinstall`, or it is not a member of the appropriate OSDBA or OSDBA for ASM groups, then create a new `oracle` user. Oracle does not support modifying an existing installation owner. See [Section 6.1.3.3, "Creating or Modifying an Oracle Software Owner User for Oracle Grid Infrastructure"](#) for a complete list of restrictions.

6.1.9.7 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 these 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.

Note: 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.

Identifying Existing User and Group IDs

To determine the user ID (`uid`) of the `grid` or `oracle` users, and the group IDs (`gid`) of the existing Oracle groups, follow these steps:

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=54321(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`.

Creating Users and Groups on the Other Cluster Nodes

To create users and groups on the other cluster nodes, repeat the following procedure on each 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, the `oinstall` and `dba` groups.

Use the `-g` option to specify the correct group ID for each group.

```
# groupadd -g 54321 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
```

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.

3. To create the `oracle` or Oracle Grid Infrastructure (`grid`) user, enter a command similar to the following:

```
# useradd -u 54321 -g oinstall -G asmadmin,asmdba grid
```

In the preceding command:

- The `-u` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- 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.

4. Set the password of the user. For example:

```
# passwd oracle
```

5. Complete user environment configuration tasks for each user as described in the section [Section 6.2, "Configuring Grid Infrastructure Software Owner User Environments"](#) on page 6-20.

6.1.10 Example of Creating Minimal Groups, Users, and Paths

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 54321 oinstall
# groupadd -g 54322 dba
# useradd -u 54321 -g oinstall -G dba oracle
# useradd -u 54322 -g oinstall -G dba grid
# mkdir -p /u01/app/12.1.0/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 ASM to administer Oracle Clusterware, Oracle ASM, and Oracle Database, and are granted SYSASM and OSOPER for 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/12.1.0/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 before installation, and 755 permissions after installation.
- /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.

6.1.11 Example of Creating Role-allocated Groups, Users, and Paths

This section contains an 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 54327 asmdba
# groupadd -g 54325 dgdba1
# groupadd -g 54335 dgdba2
# groupadd -g 54326 kmdba1
# groupadd -g 54336 kmdba2
# groupadd -g 54329 asmadmin
# groupadd -g 54328 asmoper
# useradd -u 54422 -g oinstall -G asmadmin,asmdba grid
# useradd -u 54421 -g oinstall -G dba1,backupdba1,dgdba1,kmdba1,asmdba,asmoper
oracle1
# useradd -u 54431 -g oinstall -G dba2,backupdba2,dgdba2,kmdba2,asmdba oracle2
# mkdir -p /u01/app/12.1.0/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):

- [Oracle Grid Infrastructure Groups and Users Example](#)
- [Oracle Database DB1 Groups and Users Example](#)
- [Oracle Database DB2 Groups and Users Example](#)

6.1.11.1 Oracle Grid Infrastructure Groups and Users Example

- 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).

6.1.11.2 Oracle Database DB1 Groups and Users Example

- 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.

6.1.11.3 Oracle Database DB2 Groups and Users Example

- 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.

6.2 Configuring Grid Infrastructure Software Owner User Environments

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.

This section contains the following topics:

- [Environment Requirements for Oracle Software Owners](#)
- [Procedure for Configuring Oracle Software Owner Environments](#)
- [Checking Resource Limits for the Oracle Software Installation Users](#)
- [Setting Remote Display and X11 Forwarding Configuration](#)
- [Preventing Installation Errors Caused by Terminal Output Commands](#)

6.2.1 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 software owner's environment variable `DISPLAY` environment variables 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. See [Section B.5.2, "Unset Oracle Environment Variables"](#) for more information.

6.2.2 Procedure for Configuring Oracle Software Owner Environments

To set the Oracle software owners' environments, follow these steps, for each software owner (`grid`, `oracle`):

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 `grid` user:

```
$ su - grid
```

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 installation. 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.

6.2.3 Checking Resource Limits for the Oracle Software Installation Users

For each installation software owner, check the resource limits for installation, using the following recommended ranges:

Table 6–1 Installation Owner Resource Limit Recommended Ranges

Resource Shell Limit	Resource	Soft Limit	Hard Limit
Open file descriptors	nofile	at least 1024	at least 65536
Number of processes available to a single user	nproc	at least 2047	at least 16384
Size of the stack segment of the process	stack	at least 10240 KB	at least 10240 KB, and at most 32768 KB
Maximum Locked Memory Limit	memlock	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.

6.2.4 Setting Remote Display and X11 Forwarding Configuration

If you are on a remote terminal, and the local node has only one visual (which is typical), then use the following syntax to set your user account DISPLAY environment variable:

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 `node1`, then enter the following command:

```
$ export DISPLAY=node1:0
```

To ensure that X11 forwarding does not cause the installation to fail, create a user-level SSH client configuration file for the Oracle software owner user, as follows:

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 the `~/ssh` are secured to the Grid user. For example:

```
$ ls -al .ssh
total 28
drwx----- 2 grid oinstall 4096 Jun 21 2014
drwx----- 19 grid oinstall 4096 Jun 21 2014
-rw-r--r-- 1 grid oinstall 1202 Jun 21 2014 authorized_keys
-rwx----- 1 grid oinstall 668 Jun 21 2014 id_dsa
-rwx----- 1 grid oinstall 601 Jun 21 2014 id_dsa.pub
-rwx----- 1 grid oinstall 1610 Jun 21 2014 known_hosts
```

6.2.5 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`) will cause makefile and other installation errors if they contain terminal output commands.

To avoid this problem, you must modify these 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: When SSH is not available, the Installer uses the `rsh` and `rcp` commands instead of `ssh` and `scp`.

If there are hidden files that contain `stty` commands that are loaded by the remote shell, then OUI indicates an error and stops the installation.

6.3 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.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to configure IPMI after installation

6.3.1 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.
- Each cluster member node requires an IPMI driver installed on each node.
- 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 Ethernet 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 cannot 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 operating system vendor for more information.

6.3.2 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.

Note: If you configure IPMI, and you use Grid Naming Service (GNS), then you still must configure separate addresses for the IPMI interfaces. Because the IPMI adapter is not seen directly by the host, the IPMI adapter is not visible to GNS as an address on the host.

6.3.3 Configuring the IPMI Driver

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.

6.3.3.1 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.

The following example shows how to configure the Open IPMI driver manually on Oracle Linux:

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
```

Note: 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 so that the `modprobe` commands in step 2 will be run automatically on system restart:

```
# START IPMI ON SYSTEM RESTART
/sbin/modprobe ipmi_msghandler
/sbin/modprobe ipmi_si
/sbin/modprobe ipmi_devintf
```

Note: 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 has been dynamically loaded, then the output should 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
```

In the preceding example, 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 preceding example allow the device to be accessible only by `root`. The device should only be accessed by `root`, to prevent a system vulnerability.

6.3.3.2 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:

- Enable IPMI over LAN, so that the BMC can be controlled over the management network.
- Enable dynamic IP addressing using DHCP or GNS, or configure a static IP address for the BMC.
- Establish an administrator user account and password for the BMC.
- 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.

6.3.3.2.1 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 Open IPMI Driver](#)" on page 6-26 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             : 192.0.2.10
. . .
```

- b. Turn on LAN access for the channel found. For example, where the channel is 1:

```
# ipmitool lan set 1 access on
```

4. Configure IP address settings for IPMI using one of the following procedure:

- **Using dynamic IP addressing (DHCP)**

Dynamic IP addressing is the default assumed by Oracle Universal Installer. Oracle recommends that you select this option so that nodes can be added or removed from the cluster more easily, as address settings can be assigned automatically.

Note: Use of DHCP requires a DHCP server on the subnet.

Set the channel. For example, if the channel is 1, then enter the following command to enable DHCP:

```
# ipmitool lan set 1 ipsrc dhcp
```

- **Using static IP Addressing**

If the BMC shares a network connection with the operating system, 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 lan set 1 ipaddr 192.168.0.55
# ipmitool lan set 1 netmask 255.255.255.0
# ipmitool 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 cannot 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 lan set 1 auth ADMIN MD5,PASSWORD
```

- b. List the account slots on the BMC and identify an unused slot. An unused slot that you can use is a slot less than the maximum ID, and not listed. For example, for more recent versions of the IPMI tool, you can use the command `ipmitool user summary`

```
# 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 Li
1           true   false  false   true   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  false  CALLBACK
```

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 privileges 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         : 192.0.2.10
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 : 192.0.2.1
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` and the password `mypassword`, enter the following command on `node1`:

```
$ ipmitool -H node2-ipmi -U username -P mypassword bmc info
```

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.

6.4 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 username and password when prompted during installation.

Configuring Storage for Oracle Grid Infrastructure and Oracle RAC

This chapter describes the storage configuration tasks that you must complete before you start the installer to install Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM), and that you must complete before adding an Oracle Real Application Clusters (Oracle RAC) installation to the cluster.

This chapter contains the following topics:

- [Reviewing Oracle Grid Infrastructure Storage Options](#)
- [About Shared File System Storage Configuration](#)
- [Configuring Operating System and Direct NFS Client](#)
- [Oracle Automatic Storage Management Storage Configuration](#)
- [Configuring Raw Logical Volumes on IBM: Linux on System z](#)

7.1 Reviewing Oracle Grid Infrastructure Storage Options

This section describes the supported storage options for Oracle Grid Infrastructure for a cluster, and for features running on Oracle Grid Infrastructure. It includes the following topics:

- [Supported Storage Options](#)
- [About Oracle ACFS and Oracle ADVM](#)
- [General Storage Considerations for Oracle Grid Infrastructure and Oracle RAC](#)
- [Guidelines for Using Oracle ASM Disk Groups for Storage](#)
- [Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC](#)
- [After You Have Selected Disk Storage Options](#)

See Also: The Oracle Certification site on My Oracle Support for the most current information about certified storage options:

<https://support.oracle.com>

7.1.1 Supported Storage Options

The following table shows the storage options supported for storing Oracle Clusterware and Oracle RAC files.

Note: For information about OCFS2, see the following website:

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

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.

For OCFS2 certification status, and for other cluster file system support, see the Certify page on My Oracle Support.

Table 7–1 Supported Storage Options for Oracle Clusterware and Oracle RAC

Storage Option	OCR and Voting Files	Oracle Clusterware binaries	Oracle RAC binaries	Oracle Database Files	Oracle 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 running Oracle Database on Hub Nodes for Oracle Database 11g Release 2 (11.2) and later. No for running Oracle Database on Leaf Nodes.	Yes (Oracle Database 12c Release 1 (12.1) and later)	Yes (Oracle Database 12c Release 1 (12.1) and later)
Local file system	No	Yes	Yes	No	No
OCFS2 ¹	No	No	Yes	Yes	Yes
Network file system (NFS) on a certified network-attached storage (NAS) filer Note: Direct NFS Client does not support Oracle Clusterware files.	Yes	Yes	Yes	Yes	Yes
Shared disk partitions (block devices or raw devices)	No	No	No	No	No

¹ For more information about OCFS2, see the note about OCFS2 at the beginning of this section.

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 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 Also: *Oracle Database Upgrade Guide* for information about how to prepare for upgrading an existing database

- If you do not have a storage option that provides external file redundancy, then you must configure at least three voting files locations and at least two Oracle Cluster Registry locations to provide redundancy.

Note: Oracle ACFS and OCFS2 are not supported on IBM: Linux on System z.

7.1.2 About 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). It contains the following topics:

- [About Oracle ACFS and Oracle ADVM](#)
- [Oracle ACFS and Oracle ADVM Support on Linux](#)
- [Restrictions and Guidelines for Oracle ACFS](#)

7.1.2.1 About Oracle ACFS and 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.

7.1.2.2 Oracle ACFS and Oracle ADVM Support on Linux

Oracle ACFS and Oracle ADVM are supported on Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server. [Table 7–2](#) lists the releases, platforms and kernel versions that support Oracle ACFS and Oracle ADVM. Refer to My Oracle Support Note 1369107.1 for the latest certification information about platforms and specific releases that support Oracle ACFS and Oracle ADVM.

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: Update 0, RedHat Compatible Kernel 3.10.0-123 Update 1 and later, 3.10-0-229 and later RedHat Compatible kernels ■ Unbreakable Enterprise Kernel Release: All Updates, 3.8.13-35 and later UEK 3.8.13 kernels All Updates, 4.1.12 and later UEK 4.1.12 kernels

Table 7–2 (Cont.) Platforms That Support Oracle ACFS and Oracle ADVM

Platform / Operating System	Kernel
Oracle Linux 6	<ul style="list-style-type: none"> ■ Oracle Linux 6 with Red Hat Compatible Kernel All Updates, 2.6.32-71 and later 2.6.32 RedHat Compatible kernels ■ Unbreakable Enterprise Kernel: All Updates, 2.6.39-100 and later UEK 2.6.39 kernels All Updates, 3.8.13 and later UEK 3.8.13 kernels All Updates, 4.1.12 and later UEK 4.1.12 kernels
Oracle Linux 5	<ul style="list-style-type: none"> ■ Oracle Linux 5 Update 3 with Red Hat Compatible Kernel: 2.6.18 or later ■ Unbreakable Enterprise Kernel: Update 3 and later, 2.6.39-100 and later UEK 2.6.39 kernels
Red Hat Enterprise Linux 7	<ul style="list-style-type: none"> ■ Red Hat Enterprise Linux 7: Update 0, 3.10.0-123 kernel Update 1 and later, 3.10.0-229 and later RedHat kernels
Red Hat Enterprise Linux 6	Red Hat Enterprise Linux 6
Red Hat Enterprise Linux 5	Red Hat Enterprise Linux 5 Update 3: 2.6.18 kernels
SUSE Linux Enterprise Server 11	SUSE Linux Enterprise Server 11 Service Pack 4 (SP4) SUSE Linux Enterprise Server 11 Service Pack 3 (SP3) SUSE Linux Enterprise Server 11 Service Pack 2 (SP2)

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 about platforms and specific releases that support Oracle ACFS and Oracle ADVM:
<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&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/CSP/main/article?cmd=show&type=NOT&id=854428.1>

Note: Oracle ACFS is not supported on IBM: Linux on System z.

7.1.2.3 Restrictions and Guidelines for Oracle ACFS

Note the following general restrictions and guidelines about Oracle ACFS:

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS) provides a general purpose file system. You can place Oracle Database binaries and Oracle Database files on this system, but you cannot place Oracle Clusterware files on Oracle ACFS.

For policy-managed Oracle Flex Cluster databases, be aware that 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 on Leaf Nodes.

- You cannot store Oracle Clusterware binaries and files on Oracle ACFS.
- Starting with Oracle Grid Infrastructure 12c Release 1 (12.1) for a cluster, creating Oracle data files on an Oracle ACFS file system is supported.
- You can store 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.

7.1.3 General 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.

7.1.3.1 General 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 can 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) will be restarted.

7.1.3.2 General Storage Considerations for Oracle RAC

For Standard Edition and Standard Edition 2 (SE2) Oracle RAC installations, Oracle ASM is the only 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, 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.

See Also:

- *Oracle Database 2 Day DBA* for more information about using a Fast Recovery Area
- *Oracle Automatic Storage Management Administrator's Guide* for information about failure groups and best practices for high availability and recovery

Note the following additional guidelines for supported 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.
- 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`.
- If you intend to use Oracle ASM with Oracle RAC, and 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 1 (12.1) 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.
- If you do not have a storage option that provides external file redundancy, then you must configure at least three voting file areas to provide voting file redundancy.

7.1.4 Guidelines for Using Oracle ASM Disk Groups for Storage

During Oracle Grid Infrastructure installation, you can create one disk group. 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). Note that 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.

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.

Note: The Oracle ASM instance that manages the existing disk group should be running in the Grid home.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for information about creating disk groups

7.1.5 Using Logical Volume Managers with Oracle Grid Infrastructure and Oracle RAC

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>

7.1.6 After You Have Selected Disk Storage Options

When you have determined your disk storage options, configure shared storage:

- **To use a file system**, see [Section 7.2, "About Shared File System Storage Configuration"](#).
- **To use Oracle Automatic Storage Management**, see [Section 7.4.3, "Using Disk Groups with Oracle Database Files on Oracle ASM"](#)

7.2 About Shared File System Storage Configuration

The installer suggests default locations for the Oracle Cluster Registry (OCR) or the Oracle Clusterware voting files, based on the shared storage locations detected on the server. If you choose to create these files on a file system, then review the following sections to complete storage requirements for Oracle Clusterware files:

- [Guidelines for Using a Shared File System with Oracle Grid Infrastructure](#)
- [Requirements for Oracle Grid Infrastructure Shared File System Volume Sizes](#)
- [Deciding to Use a Cluster File System for Oracle Clusterware Files](#)
- [About Direct NFS Client and Data File Storage](#)
- [Deciding to Use NFS for Data Files](#)

Note: 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.

7.2.1 Guidelines for Using a Shared File System with Oracle Grid Infrastructure

To use a shared file system for Oracle Clusterware, Oracle ASM, and Oracle RAC, the file system must comply with the following requirements:

- 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/>
- If you choose to place your Oracle Cluster Registry (OCR) files on a shared file system, then Oracle recommends that you configure your shared file systems in one of the following ways:
 - The disks used for the file system are on a highly available storage device, (for example, a RAID device).
 - At least two file systems are mounted, and use the features of Oracle Clusterware 12c Release 1 (12.1) to provide redundancy for the OCR.

- If you choose to place your database files on a shared file system, then one of the following should be true:
 - The disks used for the file system are on a highly available storage device, (for example, a RAID device).
 - The file systems consist of at least two independent file systems, with the database files on one file system, and the recovery files on a different file system.
- 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.

Note: Upgrading from Oracle9i Release 2 using the raw device or shared file for the OCR that you used for the SRVM configuration repository is not supported.

If you are upgrading Oracle Clusterware, and your existing cluster uses 100 MB OCR and 20 MB voting file partitions, then you must extend the OCR partition to at least 400 MB, and you should extend the voting file partition to 300 MB. Oracle recommends that you do not use partitions, but instead place OCR and voting files in a special type of failure group, called a quorum failure group.

All storage products must be supported by both your server and storage vendors.

See Also: *Oracle Database Quality of Service Management User's Guide* for more information about features requiring the Grid Infrastructure Management Repository

7.2.2 Requirements for Oracle Grid Infrastructure Shared File System Volume Sizes

Use [Table 7-3](#) and [Table 7-4](#) to determine the minimum size for shared file systems:

Table 7-3 Oracle Clusterware Shared File System Volume Size Minimum Requirements

File Types Stored	Number of Volumes	Volume Size
Voting files with external redundancy	1	At least 300 MB for each voting file volume
Oracle Cluster Registry (OCR) with external redundancy and the Grid Infrastructure Management Repository	1	At least 5.9 GB for the OCR volume that contains the Grid Infrastructure Management Repository (5.2 GB + 300 MB voting files + 400 MB OCR), plus 500 MB for each node for clusters greater than four nodes. For example, a six-node cluster allocation should be 6.9 GB.

Table 7–3 (Cont.) Oracle Clusterware Shared File System Volume Size Minimum

File Types Stored	Number of Volumes	Volume Size
Oracle Clusterware files (OCR and voting files) and Grid Infrastructure Management Repository with redundancy provided by Oracle software	3	<p>At least 400 MB for each OCR volume</p> <p>At least 300 MB for each voting file volume</p> <p>2 x 5.2 GB (normal redundancy):</p> <p>For 5 nodes and beyond, add 500 MB for each additional node.</p> <p>For example, for a 6 node cluster the size is 14.1 GB:</p> <ul style="list-style-type: none"> ■ Grid Infrastructure Management Repository = 2 x (5.2 GB+500 MB+500 MB) GB = 12.4 GB ■ 2 OCRs (2 x 400 MB) = 0.8 GB ■ 3 voting files (3 x 300 MB) = 0.9 GB <p>= 14.1 GB</p>

Table 7–4 Oracle RAC Shared File System Volume Size Minimum Requirements

File Types Stored	Number of Volumes	Volume Size
Oracle Database files	1	At least 1.5 GB for each volume
Recovery files	1	At least 2 GB for each volume

Note: Recovery files must be on a different volume than database files

In [Table 7–3](#) and [Table 7–4](#), the total required volume size is cumulative. For example, to store all Oracle Clusterware files on the shared file system with normal redundancy, you should have at least 2 GB of storage available over a minimum of three volumes (three separate volume locations for the OCR and two OCR mirrors, and one voting file on each volume). You should have a minimum of three physical disks, each at least 500 MB, to ensure that voting files and OCR files are on separate physical disks. If you add Oracle RAC using one volume for database files and one volume for recovery files, then you should have at least 3.5 GB available storage over two volumes, and at least 6.9 GB available total for all volumes.

Note: If you create partitions on shared partitions with `fdisk` by specifying a device size, such as `+400M`, then the actual device created may be smaller than the size requested, based on the cylinder geometry of the disk. This is due to current `fdisk` restrictions. Oracle recommends that you partition the entire disk that you allocate for use by Oracle ASM.

Note: The Grid Infrastructure Management Repository is not available on IBM: Linux on System z.

7.2.3 Deciding to Use a Cluster File System for Oracle Clusterware Files

For new installations, Oracle recommends that you use Oracle Automatic Storage Management (Oracle ASM) to store voting files and OCR files. For Linux86-64 (64-bit) platforms, Oracle provides a cluster file system, OCFS2. However, Oracle does not recommend using OCFS2 for Oracle Clusterware files.

7.2.4 About Direct NFS Client and Data File Storage

Direct NFS Client is an alternative to using kernel-managed NFS. This section contains the following information about Direct NFS Client:

- [About Direct NFS Client Storage](#)
- [About the oranfstab File and Direct NFS Client](#)
- [About Mounting NFS Storage Devices with Direct NFS Client](#)
- [Specifying Network Paths with the Oranfstab File](#)

7.2.4.1 About Direct NFS Client Storage

With Oracle Database, instead of using the operating system kernel NFS client, you can configure Oracle Database to access NFS servers directly using an Oracle internal client called Direct NFS Client. Direct NFS Client supports NFSv3, NFSv4 and NFSv4.1 protocols (excluding the Parallel NFS extension) to access the NFS server.

To enable Oracle Database to use Direct NFS Client, the NFS file systems must be mounted and available over regular NFS mounts before you start installation. Direct NFS Client manages settings after installation. If Oracle Database cannot open an NFS server using Direct NFS Client, then Oracle Database uses the platform operating system kernel NFS client. You should still set the kernel mount options as a backup, but for normal operation, Direct NFS Client uses its own NFS client.

Direct NFS Client supports up to four network paths to the NFS server. Direct NFS Client client performs load balancing across all specified paths. If a specified path fails, then Direct NFS Client reissues I/O commands over any remaining paths.

Some NFS file servers require NFS clients to connect using reserved ports. If your filer is running with reserved port checking, then you must disable reserved port checking for Direct NFS Client to operate. To disable reserved port checking, consult your NFS file server documentation.

For NFS servers that restrict port range, you can use the `insecure` option to enable clients other than `root` to connect to the NFS server. Alternatively, you can disable Direct NFS Client as described in [Section 7.3.10, "Disabling Direct NFS Client Oracle Disk Management Control of NFS"](#).

Note: Use NFS servers supported for Oracle RAC. See the following URL for support information:

<https://support.oracle.com>

7.2.4.2 About Direct NFS Client Configuration

Direct NFS Client uses either the configuration file `$ORACLE_HOME/dbs/oranfstab` or the operating system mount tab file `/etc/mtab` to find out what mount points are available. If `orafstab` is not present, then by default Direct NFS Client servers mount entries found in `/etc/mtab`. No other configuration is required. You can use `orafstab`

to specify additional specific Oracle Database operations to use Direct NFS Client. For example, you can use `oranfstab` to specify additional paths for a mount point.

Direct NFS Client supports up to four network paths to the NFS server. Direct NFS Client performs load balancing across all specified paths. If a specified path fails, then Direct NFS Client reissues I/O commands over any remaining paths.

7.2.4.3 About the `oranfstab` File and Direct NFS Client

If you use Direct NFS Client, then you can use a new file specific for Oracle data file management, `oranfstab`, to specify additional options specific for Oracle Database to Direct NFS Client. For example, you can use `oranfstab` to specify additional paths for a mount point. You can add the `oranfstab` file either to `/etc` or to `$ORACLE_HOME/dbs`.

With shared Oracle homes, when the `oranfstab` file is placed in `$ORACLE_HOME/dbs`, the entries in the file are specific to a single database. In this case, all nodes running an Oracle RAC database use the same `$ORACLE_HOME/dbs/oranfstab` file. In non-shared Oracle RAC installs, `oranfstab` must be replicated on all nodes.

When the `oranfstab` file is placed in `/etc`, then it is globally available to all Oracle databases, and can contain mount points used by all Oracle databases running on nodes in the cluster, including standalone databases. However, on Oracle RAC systems, if the `oranfstab` file is placed in `/etc`, then you must replicate the file `/etc/oranfstab` file on all nodes, and keep each `/etc/oranfstab` file synchronized on all nodes, just as you must with the `/etc/fstab` file.

See Also: [Section 7.3.1, "Configuring Operating System NFS Mount and Buffer Size Parameters"](#) for information about configuring `/etc/fstab`

In all cases, mount points must be mounted by the kernel NFS system, even when they are being served using Direct NFS Client. Refer to your vendor documentation to complete operating system NFS configuration and mounting.

Caution: Direct NFS Client cannot serve an NFS server with write size values (`wmax`) less than 32768.

7.2.4.4 About Mounting NFS Storage Devices with Direct NFS Client

Direct NFS Client determines mount point settings to NFS storage devices based on the configurations in `/etc/mtab`, which are changed with configuring the `/etc/fstab` file.

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.

Oracle Database requires that mount points be mounted by the kernel NFS system even when served through Direct NFS Client.

Note: You can have only one active Direct NFS Client implementation for each instance. Using Direct NFS Client on an instance will prevent another Direct NFS Client implementation.

If Oracle Database uses Direct NFS Client mount points configured using `oranfstab`, then it first verifies kernel NFS mounts by cross-checking entries in `oranfstab` with operating system NFS mount points. If a mismatch exists, then Direct NFS Client logs an informational message, and does not operate.

If Oracle Database cannot open an NFS server using Direct NFS Client, then Oracle Database uses the platform operating system kernel NFS client. In this case, the kernel NFS mount options must be set up as defined in "[Checking NFS Mount and Buffer Size Parameters for Oracle RAC](#)" on page 7-14. Additionally, an informational message is logged into the Oracle alert and trace files indicating that Direct NFS Client could not connect to an NFS server.

[Section 7.1.1, "Supported Storage Options"](#) lists the file types that are supported by Direct NFS Client.

The Oracle files resident on the NFS server that are served by Direct NFS Client are also accessible through the operating system kernel NFS client.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for guidelines to follow regarding managing Oracle database data files created with Direct NFS Client or kernel NFS

7.2.5 Deciding to Use NFS for Data Files

Network-attached storage (NAS) systems use NFS to access data. You can store data files on a supported NFS system.

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.

Be aware that the performance of Oracle software and databases stored on NAS devices depends on the performance of the network connection between the Oracle server and the NAS device.

For this reason, Oracle recommends that you connect the server to the NAS device using a private dedicated network connection, which should be Gigabit Ethernet or better.

7.3 Configuring Operating System and Direct NFS Client

Refer to the following sections to configure your operating system and Direct NFS Client:

- [Configuring Operating System NFS Mount and Buffer Size Parameters](#)
- [Checking Operating System NFS Mount and Buffer Size Parameters](#)
- [Checking NFS Mount and Buffer Size Parameters for Oracle RAC](#)
- [Checking TCP Network Protocol Buffer for Direct NFS Client](#)
- [Enabling Direct NFS Client Oracle Disk Manager Control of NFS](#)
- [Enabling Hybrid Columnar Compression on Direct NFS Client](#)
- [Specifying Network Paths with the Oranfstab File](#)
- [Creating Directories for Oracle Clusterware Files on Shared File Systems](#)
- [Creating Directories for Oracle Database Files on Shared File Systems](#)
- [Disabling Direct NFS Client Oracle Disk Management Control of NFS](#)

7.3.1 Configuring Operating System NFS Mount and Buffer Size Parameters

If you are using NFS for the Grid home or Oracle RAC home, then you must set up the NFS mounts on the storage to enable the following:

- The root user on the clients mounting to the storage can be considered as the root user on the file server, instead of being mapped to an anonymous user.
- The root user on the client server can create files on the NFS filesystem that are owned by root on the file server.

On NFS, you can obtain root access for clients writing to the storage by enabling `no_root_squash` on the server side. For example, to set up Oracle Clusterware file storage in the path `/vol/grid`, with nodes `node1`, `node 2`, and `node3` in the domain `mycluster.example.com`, add a line similar to the following to the `/etc/exports` file:

```
/vol/grid/ node1.mycluster.example.com(rw,no_root_squash)
node2.mycluster.example.com(rw,no_root_squash) node3.mycluster.example.com
(rw,no_root_squash)
```

If the domain or DNS is secure so that no unauthorized system can obtain an IP address on it, then you can grant root access by domain, rather than specifying particular cluster member nodes:

For example:

```
/vol/grid/ *.mycluster.example.com(rw,no_root_squash)
```

Oracle recommends that you use a secure DNS or domain, and grant root access to cluster member nodes using the domain, because using this syntax enables you to add or remove nodes without the need to reconfigure the NFS server.

If you use Grid Naming Service (GNS), then the subdomain allocated for resolution by GNS within the cluster is a secure domain. Any server without a correctly signed Grid Plug and Play (GPnP) profile cannot join the cluster, so an unauthorized system cannot obtain or use names inside the GNS subdomain.

Caution: Granting root access by domain can be used to obtain unauthorized access to systems. System administrators should see their operating system documentation for the risks associated with using `no_root_squash`.

After changing `/etc/exports`, reload the file system mount using the following command:

```
# /usr/sbin/exportfs -avr
```

7.3.2 Checking Operating System NFS Mount and Buffer Size Parameters

On Oracle Grid Infrastructure cluster member nodes, you must set the values for the NFS buffer size parameters `rsize` and `wsize` to 32768.

The NFS client-side mount options for binaries are:

```
rw,bg,hard,nointr,tcp,vers=3,timeo=600,rsize=32768,wsize=32768,actimeo=0
```

Note: The `intr` and `nointr` mount options are deprecated with Oracle Unbreakable Enterprise Linux and Oracle Linux kernels, 2.6.32 and later.

If you have Oracle Grid Infrastructure binaries on an NFS mount, then you must not include the `nosuid` option.

The NFS client-side mount options for Oracle Clusterware files (OCR and voting files) are:

```
rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,noac,vers=3,timeo=600,actimeo=0
```

Update the `/etc/fstab` file on each node with an entry containing the NFS mount options for your platform. For example, if your platform is x86-64, and you are creating a mount point for Oracle Clusterware files, then update the `/etc/fstab` files with an entry similar to the following:

```
nfs_server:/vol/grid /u02/oracle/cwfiles nfs \  
rw,bg,hard,nointr,tcp,vers=3,timeo=600,actimeo=0,rsize=32768,wsiz=32768 0 0
```

Note that mount point options are different for Oracle software binaries, Oracle Clusterware files (OCR and voting files), and data files.

To create a mount point for binaries only, provide an entry similar to the following for a binaries mount point:

```
nfs_server:/vol/bin /u02/oracle/grid nfs \  
rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,vers=3,timeo=600,actimeo=0,suid
```

See Also: My Oracle Support bulletin 359515.1, "Mount Options for Oracle Files When Used with NAS Devices" for the most current information about mount options, available from the following URL:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=359515.1>

Note: Refer to your storage vendor documentation for additional information about mount options.

7.3.3 Checking NFS Mount and Buffer Size Parameters for Oracle RAC

If you use NFS mounts for Oracle RAC files, then you must mount NFS volumes used for storing database files with special mount options on each node that has an Oracle RAC instance. When mounting an NFS file system, Oracle recommends that you use the same mount point options that your NAS vendor used when certifying the device. Refer to your device documentation or contact your vendor for information about recommended mount-point options.

Update the `/etc/fstab` file on each node with an entry similar to the following:

```
nfs_server:/vol/DATA/oradata /u02/oradata nfs\  
rw,bg,hard,nointr,tcp,vers=3,timeo=600,actimeo=0,rsize=32768,wsiz=32768 0 0
```

The mandatory mount options comprise the minimum set of mount options that you must use while mounting the NFS volumes. These mount options are essential to protect the integrity of the data and to prevent any database corruption. Failure to use

these mount options may result in the generation of file access errors. see your operating system or NAS device documentation for more information about the specific options supported on your platform.

See Also: My Oracle Support Note 359515.1 for updated NAS mount option information, available at the following URL:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=359515.1>

7.3.4 Checking TCP Network Protocol Buffer for Direct NFS Client

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
```

7.3.5 Enabling Direct NFS Client Oracle Disk Manager Control of NFS

Complete the following procedure to enable Direct NFS Client:

1. Create an `oranfstab` file with the following attributes for each NFS server you configure for 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 Direct NFS Client uses. Possible values are NFSv3, NFSv4 and NFSv4.1. The default version is NFSv3. If you select NFSv4.x, then you must configure the value in `orantstab` for `nfs_version`.
- **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 that this POSIX option 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.

See Also: *Oracle Database Performance Tuning Guide* for more information about limiting asynchronous I/O

[Example 7-1](#), [Example 7-2](#), and [Example 7-3](#) show three possible NFS server entries in `orantstab`. A single `orantstab` can have multiple NFS server entries.

2. By default, Direct NFS Client is installed in an enabled state for Oracle RAC installations. 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:
 - a. Log in as the Oracle Grid Infrastructure installation owner.
 - b. Change directory to `Grid_home/rdbms/lib`.
 - c. Enter the following commands:

```
$ make -f ins_rdbms.mk dnfs_on
```

Example 7-1 Using Local and Path NFS Server Entries

The following example uses both local and path. Because local and path are in different subnets, there is no need 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
nfs_version: nfsv3
community: private
```

Example 7-2 Using Local and Path in the Same Subnet, with dontroute

The following example shows local and path in the same subnet. `dontroute` is specified in this case:

```
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
nfs_version: nfsv4
management: 192.0.10.128

```

Example 7-3 Using Names in Place of IP Addresses, with Multiple Exports

```

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

```

7.3.6 Enabling Hybrid Columnar Compression on Direct NFS Client

To enable Hybrid Columnar Compression (HCC) on Direct NFS Client, perform the following steps:

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.

7.3.7 Specifying Network Paths with the Oranfstab File

Direct NFS Client can use up to four network paths defined in the `oranfstab` file for an NFS server. Direct NFS Client performs load balancing across all specified paths. If a specified path fails, then Direct NFS Client reissues I/O commands over any remaining paths.

Use the following SQL*Plus views for managing Direct NFS Client in a cluster environment:

- **gv\$dnfs_servers**: Shows a table of servers accessed using Direct NFS Client.
- **gv\$dnfs_files**: Shows a table of files currently open using Direct NFS Client.
- **gv\$dnfs_channels**: Shows a table of open network paths (or channels) to servers for which Direct NFS Client is providing files.
- **gv\$dnfs_stats**: Shows a table of performance statistics for Direct NFS Client.

Note: Use `v$` views for single instances, and `gv$` views for Oracle Clusterware and Oracle RAC storage.

7.3.8 Creating Directories for Oracle Clusterware Files on Shared File Systems

Use the following instructions to create directories for Oracle Clusterware files. You can also configure shared file systems for the Oracle Database and recovery files.

Note: For both NFS and OCFS2 storage, you must complete this procedure only if you want to place the Oracle Clusterware files on a separate file system from the Oracle base directory.

To create directories for the Oracle Clusterware files on separate file systems from the Oracle base directory, follow these steps:

1. If necessary, configure the shared file systems to use and mount them on each node.

Note: The mount point that you use for the file system must be identical on each node. Ensure that the file systems are configured to mount automatically when a node restarts.

2. Use the `df` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems to use. Choose a file system with a minimum of 600 MB of free disk space (one OCR and one voting file, with external redundancy).

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.

4. Note the names of the mount point directories for the file systems that you identified.
5. If the user performing installation (typically, `grid` or `oracle`) has permissions to create directories on the storage location where you plan to install Oracle Clusterware files, then OUI creates the Oracle Clusterware file directory.

If the user performing installation does not have write access, then you must create these directories manually using commands similar to the following to create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on the directory. For example, where the user is `oracle`, and the Oracle Clusterware file storage area is `cluster`:

```
# mkdir /mount_point/cluster
# chown oracle:oinstall /mount_point/cluster
# chmod 775 /mount_point/cluster
```

Note: After installation, directories in the installation path for the OCR files should be owned by `root`, and not writable by any account other than `root`.

When you have completed creating a subdirectory in the mount point directory, and set the appropriate owner, group, and permissions, you have completed OCFS2 or NFS configuration for Oracle Grid Infrastructure.

7.3.9 Creating Directories for Oracle Database Files on Shared File Systems

Use the following instructions to create directories for shared file systems for Oracle Database and recovery files (for example, for an Oracle RAC database).

1. If necessary, configure the shared file systems and mount them on each node.

Note: The mount point that you use for the file system must be identical on each node. Ensure that the file systems are configured to mount automatically when a node restarts.

2. Use the `df -h` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems:

File Type	File System Requirements
Database files	Choose either: <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 1.5 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.

4. Note the names of the mount point directories for the file systems that you identified.
5. If the user performing installation (typically, `oracle`) 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 using commands similar to the following to 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 /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```
- Recovery file directory (Fast Recovery Area):

```
# mkdir /mount_point/recovery_area
# chown oracle:oinstall /mount_point/recovery_area
# chmod 775 /mount_point/recovery_area
```

By making members of the `oinstall` group owners of these directories, this permits them to be read by multiple Oracle homes, including those with different OSDBA groups.

When you have completed creating subdirectories in each of the mount point directories, and set the appropriate owner, group, and permissions, you have completed OCFS2 or NFS configuration for Oracle Database shared storage.

7.3.10 Disabling Direct NFS Client Oracle Disk Management Control of NFS

Complete the following steps to disable Direct NFS Client:

1. Log in as the Oracle Grid Infrastructure installation owner, and disable Direct NFS Client using the following commands, where *Grid_home* is the path to the Oracle Grid Infrastructure home:

```
$ cd Grid_home/rdbms/lib
$ make -f ins_rdbms.mk dnfs_off
```

Enter these commands on each node in the cluster, or on the shared Grid home if you are using a shared home for the Oracle Grid Infrastructure installation.

2. Remove the `oranfstab` file.

Note: If you remove an NFS path that an Oracle Database is using, then you must restart the database for the change to be effective.

7.4 Oracle Automatic Storage Management Storage Configuration

Review the following sections to configure storage for Oracle Automatic Storage Management:

- [Configuring Storage for Oracle Automatic Storage Management](#)
- [About Oracle ASM with Oracle ASM Filter Driver](#)
- [Using Disk Groups with Oracle Database Files on Oracle ASM](#)
- [Configuring Oracle Automatic Storage Management Cluster File System](#)
- [Upgrading Existing Oracle ASM Instances](#)

7.4.1 Configuring Storage for Oracle Automatic Storage Management

This section describes how to configure storage for use with Oracle Automatic Storage Management.

- [Identifying Storage Requirements for Oracle Automatic Storage Management](#)
- [Creating Files on a NAS Device for Use with Oracle ASM](#)
- [Using an Existing Oracle ASM Disk Group](#)

7.4.1.1 Identifying Storage Requirements for Oracle Automatic Storage Management

To identify the storage requirements for using Oracle ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use Oracle ASM for Oracle Clusterware files (OCR and voting files), Oracle Database files, recovery files, or all files except for Oracle Clusterware or 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 Clusterware, 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. Each type of file can be stored on either Oracle ASM or a cluster file system. All the OCR files or all the voting files must use the same type of storage. You cannot have some OCR files stored in Oracle ASM and other OCR files in a cluster file system. However, you can use one type of storage for the OCR files and a different type of storage for the voting files if all files of each type use the same type of storage.
-
-

2. 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. If the voting files are in a disk group, then 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.

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:

- 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.

- 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 (two of the three disks are used by failure groups and all three disks are used by the quorum failure group) and provides three voting files and one OCR (one primary and one secondary copy). With normal redundancy, the cluster can survive the loss of one failure group.

For most installations, Oracle recommends that you select normal redundancy.

- 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 (three of the five disks are used by failure groups and all five disks are used by the quorum failure group) and provides five voting files and one OCR (one primary and two secondary copies). 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.

Note: After a disk group is created, you cannot alter the redundancy level of the disk group.

3. Determine the total amount of disk space that you require for Oracle Clusterware files, and for the database files and recovery files.

Use [Table 7-5](#) and [Table 7-6](#) 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:

Table 7-5 Oracle Clusterware Minimum Storage Space Required by Redundancy Type

Redundancy Level	Minimum Number of Disks	Oracle Cluster Registry (OCR) Files	Voting Files	Both File Types	Total Storage including Grid Infrastructure Management Repository
External	1	400 MB	300 MB	700 MB	At least 5.9 GB for a cluster with 4 nodes or less (4.5 GB + 400 MB + 300 MB). Additional space required for clusters with 5 or more nodes. For example, a six-node cluster allocation should be at least 6.9 GB: (5.2 GB + 2*(500 MB) + 400 MB + 300 MB).

Table 7–5 (Cont.) Oracle Clusterware Minimum Storage Space Required by Redundancy Type

Redundancy Level	Minimum Number of Disks	Oracle Cluster Registry (OCR) Files	Voting Files	Both File Types	Total Storage including Grid Infrastructure Management Repository
Normal	3	At least 400 MB for each failure group, or 800 MB	900 MB	1.7 GB ¹	At least 12.1 GB for a cluster with 4 nodes or less (2*5.2 GB + 2*400 MB + 3*300 MB). Additional space required for clusters with 5 or more nodes. For example, for a six-node cluster allocation should be at least 14.1 GB: (2 * (5.2 GB + 2*(500 MB)) + (2 * 400 MB) + (3 * 300 MB)).
High	5	At least 400 MB for each failure group, or 1.2 GB	1.5 GB	2.7 GB	At least 18.3 GB for a cluster with 4 nodes or less (3* 5.2 GB + 3*400 MB + 5*300 MB). Additional space required for clusters with 5 or more nodes. For example, for a six-node cluster allocation should be at least 21.3 GB: (3* (5.2 GB + 2*(500 MB)) + (3 * 400 MB) + (5 * 300 MB)).

¹ If you create a disk group during installation, then it must be at least 2 GB.

Note: If the voting files are in a disk group, be aware that disk groups with Oracle Clusterware files (OCR and voting files) have a higher minimum number of failure groups than other disk groups.

If you create a disk group as part of the installation in order to install the OCR and voting files, then the installer requires that you create these files on a disk group with at least 2 GB of available space.

Table 7–6 Total Oracle Database Storage Space Required by Redundancy Type

Redundancy Level	Minimum Number of Disks	Database Files	Recovery Files	Both File Types
External	1	1.5 GB	3 GB	4.5 GB
Normal	2	3 GB	6 GB	9 GB

Table 7–6 (Cont.) Total Oracle Database Storage Space Required by Redundancy Type

Redundancy Level	Minimum Number of Disks	Database Files	Recovery Files	Both File Types
High	3	4.5 GB	9 GB	13.5 GB

- 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. The default value is set to 1 MB.
- 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:

$$\text{total} = [2 * \text{ausize} * \text{disks}] + [\text{redundancy} * (\text{ausize} * (\text{nodes} * (\text{clients} + 1) + 30) + (64 * \text{nodes}) + 533)]$$

Where:

- redundancy = Number of mirrors: external = 1, normal = 2, high = 3.
- ausize = Metadata AU size in megabytes (default is 1 MB)
- 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 1684 MB of space:

$$[2 * 1 * 3] + [2 * (1 * (4 * (4 + 1) + 30) + (64 * 4) + 533)] = 1684 \text{ MB}$$

To ensure high availability of Oracle Clusterware files on Oracle ASM, for a normal redundancy disk group, as a general rule for most installations, you must have at least 2 GB of disk space for Oracle Clusterware files in three separate failure groups, with at least three physical disks. To ensure that the effective disk space to create Oracle Clusterware files is 2 GB, best practice suggests that you ensure at least 2.1 GB of capacity for each disk, with a total capacity of at least 6.3 GB for three disks.

- 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: Define custom failure groups after installation, using the GUI tool ASMCA, the command line tool `asmcmd`, or SQL commands.

If you define custom failure groups, then for failure groups containing database files only, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

For failure groups containing database files and clusterware files, including voting files, you must specify a minimum of three failure groups for normal redundancy disk groups, and five failure groups for high redundancy disk groups.

Disk groups containing voting files must have at least 3 failure groups for normal redundancy or at least 5 failure groups for high redundancy. Otherwise, the minimum is 2 and 3 respectively. The minimum number of failure groups applies whether or not they are custom failure groups.

7. 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. Each disk group device should 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 Grid Infrastructure and Oracle RAC require a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

Oracle recommends that if you choose to use a logical volume manager, then use the logical volume manager to represent a single LUN without striping or mirroring, so that you can minimize the impact of the additional storage layer.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for information about allocation units

7.4.1.2 Creating Files on a NAS Device for Use with Oracle ASM

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.

To create these files, follow these steps:

1. If necessary, create an exported directory for the disk group files on the NAS device.

Refer to the NAS device documentation for more information about completing this step.

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`.

See Also: ■ My Oracle Support Note 359515.1 for updated NAS mount option information, available at the following URL:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&iid=359515.1>

- For more information about editing the mount file for the operating system, see the man pages.
 - For more information about recommended mount options, see the section "[Checking NFS Mount and Buffer Size Parameters for Oracle RAC](#)".
5. Enter a command similar to the following to mount the NFS file system on the local system:

```
# mount /mnt/oracleasm
```

6. Choose a name for the disk group to create. For example: `sales1`.
7. Create a directory for the files on the NFS file system, using the disk group name as the directory name. For example:

```
# mkdir /mnt/oracleasm/nfsdg
```

8. 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/nfsdg/disk1 bs=1024k  
count=1000 oflag=direct
```

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.

9. Enter commands similar to the following to change the owner, group, and permissions on the directory and files that you created, where the installation owner is `grid`, and the OSASM group is `asmadmin`:

```
# chown -R grid:asmadmin /mnt/oracleasm  
# chmod -R 660 /mnt/oracleasm
```

10. If you plan to install Oracle RAC or a standalone Oracle Database, then during installation, edit the Oracle ASM disk discovery string to specify a regular expression that matches the file names you created. For example:

```
/mnt/oracleasm/sales1/
```

Note: During installation, disks labelled as ASMF disks or 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.

7.4.1.3 Using an Existing Oracle ASM Disk Group

Select from the following choices to store either database or recovery files in an existing Oracle ASM disk group, depending on installation method:

- If you select an installation method that runs Database Configuration Assistant in interactive mode, then you can decide whether you want to create a disk group, or to use an existing one.

The same choice is available to you if you use Database Configuration Assistant after the installation to create a database.

- If you select an installation method that runs Database Configuration Assistant in noninteractive mode, then you must choose an existing disk group for the new database; you cannot create a disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The Oracle ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine if an existing Oracle ASM disk group exists, or to determine if there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Cloud Control or the Oracle ASM command line tool (`asmcmd`) as follows:

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 listed in the previous section.

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.

7.4.2 About Oracle ASM with Oracle ASM Filter Driver

The Oracle ASM Filter Driver (Oracle ASMFDF) is installed by default with Oracle Grid Infrastructure. Oracle ASMFDF rejects write I/O requests that are not issued by Oracle software. This filter ensures that users with administrative privileges cannot inadvertently overwrite 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 ASMFDF, 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.

See Also: ■ *Oracle Automatic Storage Management Administrator's Guide* for more information about configuring storage device path persistence using Oracle ASM Filter Driver

- [Section F.4.1.4, "Deinstalling Oracle ASMLIB"](#)

7.4.3 Using Disk Groups with Oracle Database Files on Oracle ASM

Review the following sections to configure Oracle Automatic Storage Management (Oracle ASM) storage for Oracle Clusterware and Oracle Database Files:

- [Identifying and Using Existing Oracle Database Diskgroups on Oracle ASM](#)
- [Creating Diskgroups for Oracle Database Data Files](#)
- [Creating and Using Oracle ASM Credentials File](#)

7.4.3.1 Identifying and Using Existing Oracle Database Diskgroups on Oracle ASM

The following section describes how to identify existing disk groups and determine the free disk space that they contain. Optionally, identify failure groups for the Oracle ASM disk group devices. For information about Oracle ASM disk discovery, see *Oracle Automatic Storage Management Administrator's Guide*.

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.

7.4.3.2 Creating Diskgroups 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.

7.4.3.3 Creating and Using Oracle ASM Credentials File

An Oracle ASM Storage Client does not have Oracle ASM running on the nodes and uses Oracle ASM storage services in a different client cluster.

To create Oracle ASM credentials file, from the *Grid_home/bin* directory on the Storage Server, run the following command on one of the member nodes, where *credential_file* is the name and path location of the Oracle ASM credentials file you create:

```
Grid_home/bin/asmcmd mkcc client_cluster_name credential_file
```

For example:

```
Grid_home/bin/asmcmd mkcc clientcluster1 /home/grid/clientcluster1_credentials.xml
```

Copy the Oracle ASM credentials file to a secure path on the client cluster node where you run the 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 Oracle ASM credentials file. During installation, you are prompted to provide a path to the file.

-
- Note:** ■ The Oracle ASM credentials file can be used only once. If an Oracle ASM Storage Client is configured and deconfigured, you must create a new Oracle ASM credentials file.
- If the Oracle ASM credentials file is used to configure the client cluster, then it cannot be shared or reused to configure another client cluster.
-

7.4.4 Configuring Oracle Automatic Storage Management Cluster File System

Oracle ACFS is installed as part of an Oracle Grid Infrastructure installation 12c Release 1 (12.1).

You can also create a General Purpose File System configuration of ACFS using ASMCA.

See Also: [Section 7.1.2.3, "Restrictions and Guidelines for Oracle ACFS"](#) on page 7-5 for supported deployment options

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/12.1.0/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 Home**.
7. In the Create ACFS Hosted Database Home window, enter the following information:
 - **Database Home ADVM Volume Device Name:** Enter the name of the database home. The name must be unique in your enterprise. For example: `dbase_01`
 - **Database Home Mountpoint:** Enter the directory path for the mount point. For example: `/u02/acfsmounts/dbase_01`
Make a note of this mount point for future reference.
 - **Database Home Size (GB):** Enter in gigabytes the size you want the database home to be.
 - **Database Home Owner Name:** Enter the name of the Oracle Database installation owner you plan to use to install the database. For example: `oracle1`
 - **Database Home 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`
 - Click **OK** when you have completed your entries.
8. 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 the ACFS automatically in proper order when ACFS is used for an Oracle RAC 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 **Database Home Mountpoint** 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

7.4.5 Upgrading Existing Oracle ASM Instances

If you have an Oracle ASM installation from a prior release installed on your server, or in an existing Oracle Clusterware installation, then you can use Oracle Automatic Storage Management Configuration Assistant (ASMCA, located in the path `Grid_home/bin`) to upgrade the existing Oracle ASM instance to 12c Release 1 (12.1), and subsequently configure failure groups, Oracle ASM volumes and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

Note: You must first shut down all database instances and applications on the node with the existing Oracle ASM instance before upgrading it.

During installation, if you are upgrading from an Oracle ASM release before 11.2, and you chose to use Oracle ASM and ASMCA detects that there is a prior Oracle ASM version installed in another Oracle ASM home, then after installing the Oracle ASM 12c Release 1 (12.1) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an Oracle ACFS deployment by creating Oracle ASM volumes and using the upgraded Oracle ASM to create the Oracle ACFS.

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 rolling upgrade, and ASMCA is started by the root scripts during upgrade. ASMCA cannot perform a separate upgrade of Oracle ASM from a prior release to the current release.

On an existing Oracle Clusterware or Oracle RAC installation, if the prior version of Oracle ASM instances on all nodes is 11g Release 1 or later, then you are provided with the option to perform a rolling upgrade of Oracle ASM instances. If the earlier version of Oracle ASM instances on an Oracle RAC installation are from a release before 11g Release 1, then rolling upgrades cannot be performed. In that case, Oracle ASM on all nodes are upgraded to 12c Release 1 (12.1).

7.5 Configuring Raw Logical Volumes on IBM: Linux on System z

On IBM: Linux on System z, you can use raw logical volume manager (LVM) volumes for Oracle Clusterware and Automatic Storage Management files. You can create the required raw logical volumes in a volume group on either direct access storage devices (DASDs) or on SCSI devices. To configure the required raw logical volumes, follow these steps:

Note: You do not have to format FBA-type DASDs in Linux. The device name for the single whole-disk partition for FBA-type DASDs is `/dev/dasdxxxx1`.

1. If necessary, install or configure the shared DASDs that you intend to use for the disk group and restart the system.
2. Enter the following command to identify the DASDs configured on the system:

```
# more /proc/dasd/devices
```

The output from this command contains lines similar to the following:

```
0302(ECKD) at ( 94: 48) is dasdm : active at blocksize: 4096, 540000 blocks,
2109 MB
```

These lines display the following information for each DASD:

- The device number (0302)
- The device type (ECKD or FBA)
- The Linux device major and minor numbers (94: 48)
- The Linux device file name (dasdm)

In general, DASDs have device names in the form `dasdxxxx`, where `xxxx` is between one and four letters that identify the device.

- The block size and size of the device
3. From the display, identify the devices that you want to use.

If the devices displayed are FBA-type DASDs, then you do not have to configure them. You can proceed to bind them for Oracle Database files.

If you want to use ECKD-type DASDs, then enter a command similar to the following to format the DASD, if it is not already formatted:

```
# /sbin/dasdfmt -b 4096 -f /dev/dasdxxxx
```

Caution: Formatting a DASD destroys all existing data on the device. Make sure that:

- You specify the correct DASD device name
 - The DASD does not contain existing data that you want to preserve
-
-

This command formats the DASD with a block size of 4 KB and the compatible disk layout (default), which enables you to create up to three partitions on the DASD.

4. If you intend to create raw logical volumes on SCSI devices, then proceed to step 5.

If you intend to create raw logical volumes on DASDs, and you formatted the DASD with the compatible disk layout, then determine how you want to create partitions.

To create a single whole-disk partition on the device (for example, if you want to create a partition on an entire raw logical volume for database files), enter a command similar to the following:

```
# /sbin/fdasd -a /dev/dasdxxxx
```

This command creates one partition across the entire disk. You are then ready to mark devices as physical volumes. Proceed to Step 6.

To create up to three partitions on the device (for example, if you want to create partitions for individual tablespaces), enter a command similar to the following:

```
# /sbin/fdasd /dev/dasdxxxx
```

Use the following guidelines when creating partitions:

- Use the `p` command to list the partition table of the device.
- Use the `n` command to create a new partition.
- After you have created the required partitions on this device, use the `w` command to write the modified partition table to the device.
- See the `fdasd` man page for more information about creating partitions.

The partitions on a DASD have device names similar to the following, where *n* is the partition number, between 1 and 3:

```
/dev/dasdxxxxn
```

When you have completed creating partitions, you are then ready to mark devices as physical volumes. Proceed to Step 6.

5. If you intend to use SCSI devices in the volume group, then follow these steps:

- a. If necessary, install or configure the shared disk devices that you intend to use for the volume group and restart the system.
- b. To identify the device name for the disks that you want to use, enter the following command:

```
# /sbin/fdisk -l
```

SCSI devices have device names similar to the following:

```
/dev/sdxn
```

In this example, *x* is a letter that identifies the SCSI disk and *n* is the partition number. For example, `/dev/sda` is the first disk on the first SCSI bus.

- c. If necessary, use `fdisk` to create partitions on the devices that you want to use.
 - d. Use the `t` command in `fdisk` to change the system ID for the partitions that you want to use to `0x8e`.
6. Enter a command similar to the following to mark each device that you want to use in the volume group as a physical volume:

For SCSI devices:

```
# pvcreate /dev/sda1 /dev/sdb1
```

For DASD devices:

```
# pvcreate /dev/dasda1 /dev/dasdb1
```

7. To create a volume group named `oracle_vg` using the devices that you marked, enter a command similar to the following:

For SCSI devices:

```
# vgcreate oracle_vg /dev/sda1 /dev/sdb1
```

For DASD devices:

```
# vgcreate oracle_vg /dev/dasda1 /dev/dasdb1
```

8. To create the required logical volumes in the volume group that you created, enter commands similar to the following:

```
# lvcreate -L size -n lv_name vg_name
```

In this example:

- `size` is the size of the logical volume, for example `500M`
- `lv_name` is the name of the logical volume, for example `orcl_system_raw_500m`
- `vg_name` is the name of the volume group, for example `oracle_vg`

For example, to create a 500 MB logical volume for the SYSTEM tablespace for a database named `rac` in the `oracle_vg` volume group, enter the following command:

```
# lvcreate -L 500M -n rac_system_raw_500m oracle_vg
```

Note: These commands create a device name similar to the following for each logical volume:

/dev/vg_name/lv_name

9. On the other cluster nodes, enter the following commands to scan all volume groups and make them active:

```
# vgscan  
# vgchange -a y
```

Installing Oracle Grid Infrastructure for a Cluster

This chapter describes the procedures for installing Oracle Grid Infrastructure for a cluster. Oracle Grid Infrastructure consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM). If you plan afterward to install Oracle Database with Oracle Real Application Clusters (Oracle RAC), then this is phase one of a two-phase installation.

This chapter contains the following topics:

- [Installing Oracle Grid Infrastructure](#)
- [Installing Grid Infrastructure Using a Software-Only Installation](#)
- [Confirming Oracle Clusterware Function](#)
- [Confirming Oracle ASM Function for Oracle Clusterware Files](#)
- [Understanding Offline Processes in Oracle Grid Infrastructure](#)

8.1 Installing Oracle Grid Infrastructure

This section provides you with information about how to use the installer to install Oracle Grid Infrastructure. It contains the following sections:

- [Running OUI to Install Oracle Grid Infrastructure](#)
- [Installing Oracle Grid Infrastructure Using a Cluster Configuration File](#)

8.1.1 Running OUI to Install Oracle Grid Infrastructure

Complete the following steps to install Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) on your cluster. At any time during installation, if you have a question about what you are being asked to do, click the **Help** button on the OUI page.

1. On the installation media or where you have downloaded the installation binaries, run the `runInstaller` command. For example:

```
$ cd /home/grid/oracle_sw/  
$ ./runInstaller
```

2. Select one of the following installation options:

- **Install and Configure Oracle Grid Infrastructure for a Cluster**

Select this option to install either a standard cluster, or to install an Oracle Flex Cluster with Hub and Leaf Nodes.

- **Install and Configure Oracle Grid Infrastructure for a Standalone Server**

Select this option to install Oracle Grid Infrastructure in an Oracle Restart configuration. Use this option for single servers supporting Oracle Database and other applications.

See Also: *Oracle Database Installation Guide* for your platform for information about standalone server installations, as that installation option is not discussed in this document

- **Upgrade Oracle Grid Infrastructure or Oracle Automatic Storage Management**

Select this option to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle Grid Infrastructure), or to upgrade 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.

- **Install Oracle Grid Infrastructure Software Only**

Select this option to install Oracle Grid Infrastructure in a Grid home, without configuring the software.

3. Installation screens vary depending on the installation option you select. Respond to the configuration prompts as needed to configure your cluster.

Note: Click **Help** if you have any questions about the information you are asked to submit during installation.

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 be prepared to provide the SCAN names for the cluster, and the public names, and VIP names for each cluster member node.

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.

You can choose to configure the Hub and Leaf Node types manually, or you can choose to set a target size for the number of Hub Nodes in your cluster, and allow

Oracle Grid Infrastructure to maintain the number of Hub Nodes required for your cluster automatically.

When you enter the public node name, use the primary host name of each node. In other words, use the name displayed by the `hostname` command.

4. Provide information to automate root scripts, or run scripts as `root` when prompted by OUI. Click **Details** to see the log file. If `root.sh` fails on any of the nodes, then you can fix the problem and rerun `root.sh` on that node again and continue. If the problem cannot be fixed, follow the steps in [Section 10.5, "Unconfiguring Oracle Clusterware Without Removing Binaries"](#).

If you configure automation for running root scripts, and a root script fails, then you can fix the problem manually, and click **Retry** to run the root script again on nodes that failed to run the script.

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 a particular node number assignment is desired, you should run the root scripts in that order, waiting for the script to finish running on each node.

5. After `root.sh` runs on all the nodes, OUI runs Net Configuration Assistant (`netca`) and Cluster Verification Utility. These programs run without user intervention.
6. Oracle Automatic Storage Management Configuration Assistant (`asmca`) configures Oracle ASM during the installation.
7. 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`.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about using Trace File Analyzer Collector

8. 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.

When you have verified that your Oracle Grid Infrastructure installation is completed successfully, you can either use it to maintain high availability for other applications, or you can install an Oracle database.

The following is a list of additional information to note about installation:

If you are installing on Linux systems, you are using the ASM library driver (ASMLIB), and you select Oracle Automatic Storage Management (Oracle ASM) during installation, then Oracle ASM default discovery finds all disks that ASMLIB marks as Oracle ASM disks.

If you intend to install Oracle Database 12c Release 1 (12.1) with Oracle RAC, then see *Oracle Real Application Clusters Installation Guide for Linux*.

See Also: *Oracle Clusterware Administration and Deployment Guide* for cloning Oracle Grid Infrastructure, and *Oracle Real Application Clusters Administration and Deployment Guide* for information about using cloning and node addition procedures for adding Oracle RAC nodes

8.1.2 Installing Oracle Grid Infrastructure Using a Cluster Configuration File

During installation of Oracle Grid Infrastructure, you are given the option 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 OUI, which provides OUI 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.

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
.
.
.
```

Where `node-role` can have either HUB or LEAF as values.

For example:

```
mynode1 mynode1-vip /HUB
mynode2 mynode2-vip /LEAF
```

8.2 Installing Grid Infrastructure Using a Software-Only Installation

This section contains the following tasks:

- [Installing the Software Binaries](#)
- [Configuring the Software Binaries](#)
- [Configuring the Software Binaries Using a Response File](#)
- [Setting Ping Targets for Network Checks](#)

Note: Oracle recommends that only advanced users should perform the software-only installation, as this installation option requires manual postinstallation steps to enable the Oracle Grid Infrastructure software.

A software-only installation consists of installing Oracle Grid Infrastructure for a cluster on one node.

If you use the Install Grid Infrastructure Software Only option during installation, then this installs the software binaries on the local node. To complete the installation for your cluster, you must perform the additional steps of configuring Oracle

Clusterware and Oracle ASM, creating a clone of the local installation, deploying this clone on other nodes, and then adding the other nodes to the cluster.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to clone an Oracle Grid Infrastructure installation to other nodes, and then adding them to the cluster

8.2.1 Installing the Software Binaries

To perform a software-only installation:

1. Run the `runInstaller` command from the relevant directory on the Oracle Database 12c Release 1 (12.1) installation media or download directory. For example:

```
$ cd /home/grid/oracle_sw
$ ./runInstaller
```

2. Complete a software-only installation of Oracle Grid Infrastructure on the first node.
3. When the software has been installed, run the `oraInstRoot.sh` script when prompted.
4. 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.

However, ignore the instruction to run the `roothas.pl` script, unless you intend to install Oracle Grid Infrastructure on a standalone server (Oracle Restart).

5. Verify that all of the cluster nodes meet the installation requirements using the command `runcluvfy.sh stage -pre crsinst -n node_list`. Ensure that you have completed all storage and server preinstallation requirements.
6. Use Oracle Universal Installer as described in steps 1 through 4 to install the Oracle Grid Infrastructure software on every remaining node that you want to include in the cluster, and complete a software-only installation of Oracle Grid Infrastructure on every node.
7. Configure the cluster using the full OUI configuration wizard GUI as described in [Section 8.2.2, "Configuring the Software Binaries,"](#) or configure the cluster using a response file as described in section [Section 8.2.3, "Configuring the Software Binaries Using a Response File."](#)

8.2.2 Configuring the Software Binaries

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode:

1. Log in to a terminal as the Grid infrastructure installation owner, and change directory to `Grid_home/crs/config`.
2. Enter the following command:

```
$ ./config.sh
```

The configuration script starts OUI in Configuration Wizard mode. Provide information as needed for configuration. 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.

3. When you complete inputs, OUI shows you the Summary page, listing all inputs 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.

4. When prompted, run root scripts.
5. When you confirm that all root scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.

8.2.3 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. This section provides the procedure for completing configuration after the software is installed or copied on nodes, using the configuration wizard utility (`config.sh`).

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about the configuration wizard.

To configure the Oracle Grid Infrastructure software binaries using a response file:

1. As the Oracle Grid Infrastructure installation owner (`grid`) start OUI 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/crs/config/config.sh [-debug] [-silent -responseFile filename]
```

For example:

```
$ cd /u01/app/12.1.0/grid/crs/config/  
$ ./config.sh -responseFile /u01/app/grid/response/response_file.rsp
```

The configuration script starts OUI 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 inputs, OUI shows you the Summary page, listing all inputs 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.

8.2.4 Setting Ping Targets for Network Checks

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, using the installer as follows:

```
./runInstaller 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 them respond, then the network is considered to be offline. Addresses outside the cluster, such as switch or router address, should be used.

For example:

```
./runInstaller oracle_install_crs_Ping_Targets=192.0.2.1,192.0.2.2
```

8.3 Confirming Oracle Clusterware Function

After installation, log in as root, and use the following command syntax on each node to confirm that your Oracle Clusterware installation is installed and running correctly:

```
crsctl check cluster
```

For example:

```
$ crsctl check cluster
```

```
CRS-4537 Cluster Ready Services is online  
CRS-4529 Cluster Synchronization Services is online  
CRS-4533 Event Manager is online
```

Caution: 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.

8.4 Confirming Oracle ASM Function for Oracle Clusterware Files

If you installed the OCR and voting files on Oracle ASM, then use the following command syntax as the Oracle Grid Infrastructure installation owner to confirm that your Oracle ASM installation is running:

```
srvctl status asm
```

For example:

```
$ srvctl status asm  
ASM is running on node1,node2
```

Oracle ASM is running only if it is needed for Oracle Clusterware files. If you have not installed OCR and voting files on Oracle ASM, then the Oracle ASM instance should be down.

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.

8.5 Understanding Offline Processes in Oracle Grid Infrastructure

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 will prompt you to activate the required offline resource.

Oracle Grid Infrastructure Postinstallation Procedures

This chapter describes how to complete the postinstallation tasks after you have installed the Oracle Grid Infrastructure software.

This chapter contains the following topics:

- [Required Postinstallation Tasks](#)
- [Recommended Postinstallation Tasks](#)
- [Using Earlier Oracle Database Releases with Oracle Grid Infrastructure](#)
- [Modifying Oracle Clusterware Binaries After Installation](#)

9.1 Required Postinstallation Tasks

Download and install patch updates. See the My Oracle Support website for required patch updates for your installation.

To download required patch updates:

1. Use a Web browser to view the My Oracle Support website:

<https://support.oracle.com>

2. Log in to My Oracle Support.

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. On the Patches and Updates page, click **Product or Family (Advanced)**.
5. In the Product field, select **Oracle Database**.
6. In the Release field, select one or more release numbers. For example, Oracle 12.1.0.1.0.
7. Click **Search**.
8. Any available patch updates are displayed in the **Patch Search** page.
9. Click the patch number to download the patch.
10. Select the patch number and click **Read Me**. The README page contains information about the patch set and how to apply the patches to your installation.

11. Return to the Patch Set page, click **Download**, and save the file on your system.
12. Use the unzip utility provided with Oracle Database 12c Release 1 (12.1) to uncompress the Oracle patch updates that you downloaded from My Oracle Support. The unzip utility is located in the \$ORACLE_HOME/bin directory.
13. See [Appendix B, "How to Upgrade to Oracle Grid Infrastructure 12c Release 1"](#) for information about how to stop database processes in preparation for installing patches.

9.2 Recommended Postinstallation Tasks

Oracle recommends that you complete the following tasks as needed after installing Oracle Grid Infrastructure:

- [Tuning Semaphore Parameters](#)
- [Create a Fast Recovery Area Disk Group](#)
- [Checking the SCAN Configuration](#)
- [Downloading and Installing the ORAchk Health Check Tool](#)
- [Setting Resource Limits for Oracle Clusterware and Associated Databases and Applications](#)

9.2.1 Tuning Semaphore Parameters

Use the following guidelines only 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 `semmsl` (semaphores for each set) to 250.
4. Set `semnli` (total semaphores sets) to `semms` divided by `semmsl`, rounded up to the nearest multiple of 1024.

See Also: My Oracle Support note 226209.01, "Linux: How to Check Current Shared Memory, Semaphore Values," which is available from the following URL:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=226209.1>

9.2.2 Create a Fast Recovery Area Disk Group

During installation, by default you can create one disk group. 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.

9.2.2.1 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. 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. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

When you enable Fast Recovery in the `init.ora` file, all RMAN backups, archive logs, control file automatic backups, and database copies are written to the Fast Recovery Area. RMAN automatically manages files in the Fast Recovery Area by deleting obsolete backups and archive 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 Area files in the same disk group. However, Oracle recommends that you create a separate Fast Recovery Area disk group to reduce storage device contention.

The Fast Recovery Area is enabled by setting `DB_RECOVERY_FILE_DEST`. 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 should be 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 one fast recovery area disk group on disks with 150 gigabyte (GB) of storage, shared by three different databases. You can set the size of the fast recovery area for each database depending on the importance of each database. For example, if `test1` is your least important database, `products` is of greater importance and `orders` 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 `test1`, 50 GB for `products`, and 70 GB for `orders`.

See Also: *Oracle Enterprise Manager Real Application Clusters Guide Online Help*

9.2.2.2 Creating the Fast Recovery Area Disk Group

To create a Fast Recovery Area disk group:

1. Navigate to the Grid home bin directory, and start Oracle ASM Configuration Assistant (ASMCA). For example:

```
$ cd /u01/app/12.1.0/grid/bin
$ ./asmca
```

2. ASMCA opens at the Disk Groups tab. Click **Create** to create a new disk group.
3. The Create Disk Groups window opens.

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.

In the Select Member Disks field, select eligible disks to be added to the Fast Recovery Area, and click **OK**.

4. The Diskgroup Creation window opens to inform you when disk group creation is complete. Click **OK**.
5. Click **Exit**.

9.2.3 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. 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 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 scan
```

```
Checking Single Client Access Name (SCAN)...
```

```
Checking TCP connectivity to SCAN Listeners...
```

```
TCP connectivity to SCAN Listeners exists on all cluster nodes
```

```
Checking name resolution setup for "node1.example.com"...
```

```
Verification of SCAN VIP and Listener setup passed
```

```
Verification of scan was successful.
```

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

9.2.4 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, 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
- E-Business Suite

For information about configuring and running the ORAchk utility, refer to My Oracle Support note 1268927.1:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1268927.1>

9.2.5 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. These 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 `CRS_LIMIT_NPROC` parameter to a high value.

9.3 Using Earlier Oracle Database Releases with Oracle Grid Infrastructure

Review the following sections for information about using earlier Oracle Database releases with Oracle Grid Infrastructure 12c Release 1 (12.1) installations:

- [General Restrictions for Using Earlier Oracle Database Versions](#)
- [Managing Server Pools with Earlier Database Versions](#)
- [Using ASMCA to Administer Disk Groups for Earlier Database Versions](#)
- [Making Oracle ASM Available to Earlier Oracle Database Releases](#)
- [Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x](#)
- [Using the Correct LSNRCTL Commands](#)

9.3.1 General Restrictions for Using Earlier Oracle Database Versions

You can use Oracle Database 10g Release 2 and Oracle Database 11g Release 1 and 2 with Oracle Clusterware 12c Release 1 (12.1).

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 higher. Earlier releases of Oracle Database cannot be installed on Oracle ACFS because these releases were not designed to use Oracle ACFS.

If you upgrade an existing version of Oracle Clusterware and Oracle ASM to Oracle Grid Infrastructure 11g or later (which includes Oracle Clusterware and Oracle ASM), and you also plan to upgrade your Oracle RAC database to 12c Release 1 (12.1), then the required configuration of existing databases is completed automatically when you complete the Oracle RAC upgrade, and this section does not concern you.

Note: Before you start an Oracle RAC or Oracle Database installation on an Oracle Clusterware 12c Release 1 (12.1) installation, if you are upgrading from Oracle Database 11g Release 1 (11.1.0.7 or 11.1.0.6), or Oracle Database 10g Release 2 (10.2.0.4), then Oracle recommends that you check for the latest recommended patches for the release you are upgrading from, and install those patches as needed on your existing database installations before upgrading.

For more information on recommended patches, see Oracle 12c Upgrade Companion (My Oracle Support Note 1462240.1):

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1462240.1>

9.3.2 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 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 12.1 Oracle Grid Infrastructure binaries directory in the Grid home. For example:

```
# cd /u01/app/12.1.0/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/11.2.0/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

9.3.3 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 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.

If you have Oracle Database 10g Release 2 databases that use Oracle ASM for storage, then set the `SQLNET.ALLOWED_LOGON_VERSION=8` in the `$crs_home/network/admin/sqlnet.ora` file.

9.3.4 Using ASMCA to Administer Disk Groups for Earlier Database Versions

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 (11.2), 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 (12.1)

9.3.5 Pinning Cluster Nodes for Oracle Database Release 10.x or 11.x

When Oracle Clusterware 12c Release 1 (12.1) is installed on a cluster with no previous Oracle software version, it configures the cluster nodes dynamically, which is compatible with Oracle Database Release 11.2 and later, but Oracle Database 10g and 11.1 require a persistent configuration. This process of association of a node name with a node number is called pinning.

Note: During an upgrade, all cluster member nodes are pinned automatically, and no manual pinning is required for existing databases. This procedure is required only if you install earlier database versions after installing Oracle Grid Infrastructure 12c Release 1 (12.1) software.

To pin a node in preparation for installing or using an earlier Oracle Database version, use `Grid_home/bin/crsctl` with the following command syntax, where `nodes` is a space-delimited list of one or more nodes in the cluster whose configuration you want to pin:

```
crsctl pin css -n nodes
```

For example, to pin nodes `node3` and `node4`, log in as `root` and enter the following command:

```
$ crsctl pin css -n node3 node4
```

To determine if a node is in a pinned or unpinned state, use `Grid_home/bin/olsnodes` with the following command syntax:

To list all pinned nodes:

```
olsnodes -t -n
```

For example:

```
# /u01/app/12.1.0/grid/bin/olsnodes -t -n
node1 1      Pinned
node2 2      Pinned
node3 3      Pinned
node4 4      Pinned
```

To list the state of a particular node:

```
olsnodes -t -n node3
```

For example:

```
# /u01/app/12.1.0/grid/bin/olsnodes -t -n node3
node3 3      Pinned
```

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about pinning and unpinning nodes

9.3.6 Using the Correct LSNRCTL Commands

To administer 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.

9.4 Modifying Oracle Clusterware Binaries After Installation

After installation, if you need to modify the Oracle Clusterware configuration, then you must unlock the Grid home.

For example, if you want to apply a one-off patch, or if you want to modify an Oracle Clusterware configuration to run IPC traffic over RDS on the interconnect instead of using the default UDP, then you must unlock the Grid home.

Note: Before relinking executables, you must shut down all executables that run in the Oracle home directory that you are unlocking and relinking. In addition, shut down applications linked with Oracle shared libraries.

Unlock the home using the following procedure:

1. Log in as `root`, and 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 -crshome Grid_home`, where `Grid_home` is the path to your Grid infrastructure home. For example, with the Grid home `/u01/app/12.1.0/grid`, enter the following command:

```
# cd /u01/app/12.1.0/grid/crs/install
```

```
# perl rootcrs.sh -unlock -crshome /u01/app/12.1.0/grid
```

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 Oracle Grid Infrastructure installation owner and run the command `Grid_home/bin/relink`.

3. Relock the Grid home and restart the cluster using the following command:

```
# perl rootcrs.sh -patch
```

4. 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`".

How to Modify or Deinstall Oracle Grid Infrastructure

This chapter describes how to modify or remove Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).

Oracle recommends that you use the deinstallation tool to remove the entire Oracle home associated with the Oracle Database, Oracle Clusterware, Oracle ASM, Oracle RAC, or Oracle Database client installation. Oracle does not support the removal of individual products or components.

This chapter contains the following topics:

- [Deciding When to Deinstall Oracle Clusterware](#)
- [Migrating Standalone Grid Infrastructure Servers to a Cluster](#)
- [Relinking Oracle Grid Infrastructure for a Cluster Binaries](#)
- [Changing the Oracle Grid Infrastructure Home Path](#)
- [Unconfiguring Oracle Clusterware Without Removing Binaries](#)
- [Removing Oracle Clusterware and Oracle ASM](#)

See Also: Product-specific documentation for requirements and restrictions to remove an individual product

10.1 Deciding When to Deinstall Oracle Clusterware

Remove installed components in the following situations:

- You have successfully installed Oracle Clusterware, and you want to remove the Oracle Clusterware installation, either in an educational environment, or a test environment.
- You have encountered errors during or after installing or upgrading Oracle Clusterware, and you want to reattempt an installation.
- Your installation or upgrade stopped because of a hardware or operating system failure.
- You are advised by Oracle Support to reinstall Oracle Clusterware.

10.2 Migrating Standalone 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:

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about the configuration wizard.

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 listeners:

```
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. Log in as root, and change directory to `Grid home/crs/install`. For example:


```
# cd /u01/app/12.1.0/grid/crs/install
```
3. Stop all of the databases, services, and listeners that you discovered in step 1.
4. If present, unmount all Oracle Automatic Storage Management Cluster File System (Oracle ACFS) filesystems.
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, choose if you want to 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. Use the deinstallation tool to remove the Oracle Restart software, but with all disk groups intact.
- b. Proceed to step 7.

Installing in a Different Location than Oracle Restart

- a. Install Oracle Grid Infrastructure for a cluster in the new Grid home software location.
7. As the Oracle Grid Infrastructure installation owner, run Oracle Clusterware Configuration Wizard, and save and stage the response file. For example:


```
$ Grid_home/crs/config/config.sh -silent -responseFile $HOME/GI.rsp
```
 8. Run `root.sh` for the Oracle Clusterware Configuration Wizard.
 9. Mount the Oracle ASM disk group used by Oracle Restart.
 10. 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.

11. 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/11.2.0/grid/bin/srvctl add filesystem -device
/dev/asm/db1 -diskgroup ORestartData -volume db1 -mountpointpath
/u01/app/grid/product/11.2.0/db1 -user grid
```

12. 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 -oraclehome $ORACLE_HOME -node
nodename
```

For example, first verify that the `ORACLE_HOME` environment variable is set to the location of the database home directory.

Next, to add the database name `mydb`, and the service `myservice`, enter the following commands:

```
srvctl add database -db mydb -oraclehome $ORACLE_HOME -node node1
```

13. Add each service to the database, using the command `srvctl add service`. For example:

```
srvctl add service -db mydb -service myservice
```

10.3 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 -patch
```

You must relink the Oracle Clusterware and Oracle ASM binaries every time you apply an operating system patch or after an operating system upgrade.

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.

10.4 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 enter the command `crsctl stop crs`. For example:

```
$ cd /u01/app/12.1.0/grid/bin
$ ./crsctl stop crs
```
3. Detach the existing Grid home by running the following command, where `/u01/app/12.1.0/grid` is the existing Grid home location:

```
$ /u01/app/12.1.0/grid/oui/bin/runInstaller -silent -waitforcompletion\
-detachHome ORACLE_HOME='/u01/app/12.1.0/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/12.1.0/grid` and the new Grid home is `/u01/app/12c/`:

```
# mkdir /u01/app/12c
# mv /u01/app/12.1.0/grid /u01/app/12c
```
5. 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.
6. As `root` again, enter the following command to start up in the new home location:

```
# cd /u01/app/12c/crs/install
# rootcrs.sh -patch -dstcrshome /u01/app/12c/
```
7. Repeat steps 1 through 4 on each cluster member node.

You must relink the Oracle Clusterware and Oracle ASM binaries every time you move the Grid home.

10.5 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/12.1.0/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.

5. After deconfiguring an Oracle ASM Storage Client, run the following command on the Storage Server:

```
asmcmd rmcc client_cluster_name
```

10.6 Removing Oracle Clusterware and Oracle ASM

The `deinstall` command removes Oracle Clusterware and Oracle ASM from your server. The following sections describe the deinstallation tool, and provide information about additional options to use the deinstallation tool:

- [About the Deinstallation Tool](#)
- [Deinstallation Tool Command Example for Oracle Grid Infrastructure](#)
- [Deinstallation Response File Example for Grid Infrastructure for a Cluster](#)

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 12.1.0.1 installation media to remove Oracle software from an existing 11.2.0.4 Oracle home.

10.6.1 About the Deinstallation Tool

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, Oracle Database Client, or Oracle Grid Infrastructure 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 uses the information you provide, plus information gathered from the software home to create a response file. You can alternatively supply a response file generated previously by the `deinstall` command using the `-checkonly` option, or by editing the response file template.

The deinstallation tool stops Oracle software, and removes Oracle software and configuration files on the operating system for a specific Oracle home. If you run the deinstallation tool to remove Oracle Grid Infrastructure, then the deinstaller prompts you to run the `rootcrs.sh` script, as the `root` user, to deconfigure Oracle Grid Infrastructure or `roothas.sh` script to deconfigure Oracle Grid Infrastructure for standalone server.

Note: 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.

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.

Caution: 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 deinstallation tool removes the following files and directory contents in the Oracle base directory of the Oracle Database installation owner:

- admin
- cfgtoollogs
- checkpoints
- diag
- oradata
- flash_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.

In addition, 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) volumes.
 - If Grid Naming Service (GNS) is in use, then your DNS administrator must delete the entry for the subdomain from DNS.
-

The default method for running the deinstallation tool is from the deinstall directory in the Oracle home as the installation owner:

```
$ $ORACLE_HOME/deinstall/deinstall
```

The deinstall command uses the following syntax, where variable content is indicated in italics:

```
deinstall [-silent] [-checkonly] [-local] [-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] [-logdir complete path of log directory] [-help]
```

To run the deinstallation tool from the database installation media, use the runInstaller command with the -deinstall option, followed by the -home option to specify the path of the Oracle home you want to remove using the following syntax, where variable content is indicated in italics:

```
runInstaller -deinstall -home complete path of Oracle home [-silent] [-checkonly]
[-local] [-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] [-logdir complete path of log directory]
[-help]
```

Provide information about your servers as prompted or accept the defaults.

The deinstallation tool stops Oracle software, and removes Oracle software and configuration files on the operating system.

In addition, you can run the deinstallation tool with a response file, or select the following options to run the tool:

- `-home`

Use this flag to indicate the home path of the Oracle home to check or deinstall.

If you run `deinstall` from the `$ORACLE_HOME/deinstall` path, then the `-home` flag is not required because the tool identifies the location of the home where it is run. If you use `runInstaller -deinstall` from the installation media, then `-home` is mandatory.

To deinstall Oracle software using the `deinstall` command in the Oracle home you plan to deinstall, provide a parameter file located outside the Oracle home, and do not use the `-home` flag.

- `-silent`

Use this flag to run the deinstallation tool in noninteractive mode.

- A working system that it can access to determine the installation and configuration information. The `-silent` 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.

You can generate a response file to use or modify by running the tool with the `-checkonly` 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 `-silent` flag. The `-silent` flag does not work with failed installations

- `-checkonly`

Use this flag to check the status of the Oracle software home configuration. Running the `deinstall` command with the `-checkonly` flag does not remove the Oracle configuration. The `-checkonly` flag generates a response file that you can use with the `deinstall` command and `-silent` option.

- `-local`

Use this flag on a multinode environment to deinstall Oracle software in a cluster.

When you run `deinstall` with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where `deinstall` is run). It does not deinstall or deconfigure Oracle software on remote nodes.

- `-paramfile complete path of input response file`

Use this flag to run `deinstall` 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.

The default location of the response file depends on the location of `deinstall`:

- From the installation media or stage location: `stagelocation/response`
where `stagelocation` is the path of the base directory in the installation media, or in the staged files location.
- After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`

- `-params [name1=value name2=value name3=value ...]`
Use this flag with a response file to override one or more values to change in a response file you have created.
- `-o complete path of directory for saving response files`
Use this flag to provide a path other than the default location where the response file (`deinstall.rsp.tmpl`) is saved.
The default location of the response file depends on the location of `deinstall`:
 - From the installation media or stage location: `stagelocation/response`
where `stagelocation` is the path of the base directory in the installation media, or in the staged files location.
 - After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`
- `-tmpdir complete path of temporary directory`
Use this flag to specify a non-default location where the deinstallation tool writes the temporary files for the deinstallation.
- `-logdir complete path of log directory`
Use this flag to specify a non-default location where the deinstallation tool writes the log files for the deinstallation.
- `-help`
Use the help option (`-help`) to get additional information about the deinstallation tool option flags.

10.6.1.1 Deinstalling Previous Release Grid Home

For upgrades from previous releases, if you want to deinstall the previous release Grid home, then as the root user, you must manually change the permissions of the previous release Grid home, and then run the deinstallation tool.

For example:

```
# chown -R grid:oinstall /u01/app/grid/11.2.0
# chmod -R 775 /u01/app/grid/11.2.0
```

In this example, `/u01/app/grid/11.2.0` is the previous release Grid home.

10.6.2 Deinstallation Tool Command Example for Oracle Grid Infrastructure

If you run the deinstallation tool using the deinstallation tool from the `$ORACLE_HOME/deinstall` folder, then the deinstallation starts without prompting you for an `ORACLE_HOME`.

Use the optional flag `-paramfile` to provide a path to a response file.

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/12.1.0/grid/` is the path to the Grid home that you want to remove:

```
$ cd /directory_path/
$ ./runInstaller -deinstall -home /u01/app/12.1.0/grid
```

The following example uses a response file in the software owner location `/home/usr/grid`:

```
$ cd /directory_path/runInstaller
$ ./runInstaller -deinstall -paramfile /home/usr/grid/my_db_paramfile.tmpl
```

10.6.3 Deinstallation Response File Example for 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/12.1.0/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, 2006 Oracle Corporation. All rights reserved.
#Mon Feb 17 00:08:58 PST 2014
LOCAL_NODE=node1
HOME_TYPE=CRS
ASM_REDUNDANCY=\
ORACLE_BASE=/u01/app/12.1.0/grid/
VIP1_MASK=255.255.252.0
VOTING_DISKS=/u02/storage/grid/vdsk
SCAN_PORT=1522
silent=true
ASM_UPGRADE=false
ORA_CRS_HOME=/u01/app/12.1.0/grid
GPNPCONFIGDIR=$ORACLE_HOME
LOGDIR=/home/grid/SH/deinstall/logs/
GPNPGCONFIGDIR=$ORACLE_HOME
ORACLE_OWNER=grid
NODELIST=node1,node2
CRS_STORAGE_OPTION=2
NETWORKS="eth0"/192.0.2.1\:\:public,"eth1"/10.0.0.1\:\:cluster_interconnect
VIP1_IP=192.0.2.2
NETCFGJAR_NAME=netcfg.jar
ORA_DBA_GROUP=dba
CLUSTER_NODES=node1,node2
JREDIR=/u01/app/12.1.0/grid/jdk/jre
VIP1_IF=eth0
REMOTE_NODES=node2
VIP2_MASK=255.255.252.0
ORA_ASM_GROUP=asm
LANGUAGE_ID=AMERICAN_AMERICA.WE8ISO8859P1
CSS_LEASEDURATION=400
NODE_NAME_LIST=node1,node2
SCAN_NAME=node1scn
SHAREJAR_NAME=share.jar
HELPJAR_NAME=help4.jar
SILENT=false
local=false
INVENTORY_LOCATION=/u01/app/oraInventory
GNS_CONF=false
JEWJAR_NAME=jewt4.jar
OCR_LOCATIONS=/u02/storage/grid/ocr
EMBASEJAR_NAME=oemlt.jar
ORACLE_HOME=/u01/app/12.1.0/grid
CRS_HOME=true
```

```
VIP2_IP=192.0.2.4
ASM_IN_HOME=n
EWTJAR_NAME=ewt3.jar
HOST_NAME_LIST=node1,node2
JLIBDIR=/u01/app/12.1.0/grid/jlib
VIP2_IF=eth0
VNDR_CLUSTER=false
CRS_NODEVIPS='node1-vip/255.255.252.0/eth0,node2-vip/255.255.252.0/eth0'
CLUSTER_NAME=node1-cluster
```

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 "eth2"/192.0.2.2:vip1_ip
```

Troubleshooting the Oracle Grid Infrastructure Installation Process

This appendix provides troubleshooting information for installing Oracle Grid Infrastructure.

See Also: The Oracle Database 12c Release 1 (12.1) Oracle RAC documentation set in the Documentation directory:

- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*

This appendix contains the following topics:

- [Best Practices for Contacting Oracle Support](#)
- [General Installation Issues](#)
- [Interpreting CVU "Unknown" Output Messages Using Verbose Mode](#)
- [Interpreting CVU Messages About Oracle Grid Infrastructure Setup](#)
- [About the Oracle Clusterware Alert Log](#)
- [Missing Operating System Packages On Linux](#)
- [Performing Cluster Diagnostics During Oracle Grid Infrastructure Installations](#)
- [About Using CVU Cluster Healthchecks After Installation](#)
- [Interconnect Configuration Issues](#)
- [SCAN VIP and SCAN Listener Issues](#)
- [Storage Configuration Issues](#)
- [Failed or Incomplete Installations and Upgrades](#)

A.1 Best Practices for Contacting Oracle Support

If you find that it is necessary for you to contact Oracle Support to report an issue, then Oracle recommends that you follow these guidelines when you enter your service request:

- Provide a clear explanation of the problem, including exact error messages.
- Provide an explanation of any steps you have taken to troubleshoot issues, and the results of these steps.

- Provide exact versions (major release and patch release) of the affected software.
- Provide a step-by-step procedure of what actions you carried out when you encountered the problem, so that Oracle Support can reproduce the problem.
- Provide an evaluation of the effect of the issue, including affected deadlines and costs.
- Provide screen shots, logs, Remote Diagnostic Agent (RDA) output, or other relevant information.

A.2 General Installation Issues

The following is a list of examples of types of errors that can occur during installation. It contains the following issues:

- `root.sh` failed to complete with error messages such as: Start of resource "ora.cluster_interconnect.haip" failed...
- During installation of Oracle Clusterware, check for Oracle ASM disks failed with the error: PRVF-5150 : Path /dev/mapper/<alias> is not a valid path on all nodes
- An error occurred while trying to get the disks
- CRS-5018:(:CLSN00037:) Removed unused HAIP route:
- Could not execute auto check for display colors using command /usr/X11R6/bin/xdpyinfo
- CRS-5823:Could not initialize agent framework
- Failed to connect to server, Connection refused by server, or Can't open display
- Failed to initialize ocrconfig
- INS-32026 INSTALL_COMMON_HINT_DATABASE_LOCATION_ERROR
- CLSRSC-444: Run `root.sh` command on the Node with OUI session
- MEMORY_TARGET not supported on this system
- Nodes unavailable for selection from the OUI Node Selection screen
- Node nodename is unreachable
- PROT-8: Failed to import data from specified file to the cluster registry
- PRVE-0038 : The SSH LoginGraceTime setting, or fatal: Timeout before authentication
- Timed out waiting for the CRS stack to start
- YPBINDPROC_DOMAIN: Domain not bound

root.sh failed to complete with error messages such as: Start of resource "ora.cluster_interconnect.haip" failed...

Cause: When configuring public and private network interfaces for Oracle RAC, you must enable ARP. Highly Available IP (HAIP) addresses do not require ARP on the public network, but for VIP failover, you will need to enable ARP. Do not configure NOARP.

Action: Configure the `hsi0` (or `eth`) device to use ARP protocol by running the following command:

```
# ifconfig hsi0 arp
```

During installation of Oracle Clusterware, check for Oracle ASM disks failed with the error: PRVF-5150 : Path /dev/mapper/<alias> is not a valid path on all nodes

Cause: This error may occur on Red Hat Enterprise Linux 6.3 because read access to the `/etc/multipath.conf` file is not set.

Action: To resolve this, as root, add `+r` to `/etc/multipath.conf` as follows:

```
#chmod +r /etc/multipath.conf
```

An error occurred while trying to get the disks

Cause: There is an entry in `/etc/oratab` pointing to a non-existent Oracle home. The OUI log file should show the following error: "java.io.IOException: /home/oracle/OraHome/bin/kfod: not found"

Action: Remove the entry in `/etc/oratab` pointing to a non-existing Oracle home.

CRS-5018:(CLSN00037:) Removed unused HAIP route:

Cause: Typically, this error indicates that something (usually Zero Configuration Networking, `zeroconf`) has created the indicated route that is conflicting with the HAIP code. The error indicates that the Oracle software has removed the route to ensure appropriate stack functioning.

Action: The Zero Configuration Networking feature should be disabled when using Oracle Clusterware, as this feature may cause communication issues between cluster member nodes.

To disable Zero Configuration Networking:

1. Log in as root.
2. Change directory to `/etc/sysconfig`.
3. Create a copy of `/etc/sysconfig/network`. For example:

```
# cp network network_old
```

4. Use a text editor to open the file `/etc/sysconfig/network`.
5. Check the file for the value for `NOZEROCONF` to confirm that it is set to `yes`. If you do not find this parameter in the file, then append the following entry to the file:

```
NOZEROCONF=yes
```

Save the file after you update this setting.

6. Restart the network services. For example:

```
# service network restart
```

7. Repeat this process on each cluster member node.

Could not execute auto check for display colors using command /usr/X11R6/bin/xdpyinfo

Cause: Either the `DISPLAY` variable is not set, or the user running the installation is not authorized to open an X window. This can occur if you run the installation from a remote terminal, or if you use an `su` command to change from a user that is authorized to open an X window to a user account that is not authorized to open an X window on the display, such as a lower-privileged user opening windows on the `root` user's console display.

Action: Run the command `echo $DISPLAY` to ensure that the variable is set to the correct visual or to the correct host. If the `display` variable is set correctly then

either ensure that you are logged in as the user authorized to open an X window, or run the command `xhost +` to allow any user to open an X window.

If you are logged in locally on the server console as `root`, and used the `su -` command to change to the Oracle Grid Infrastructure installation owner, then log out of the server, and log back in as the grid installation owner.

CRS-5823:Could not initialize agent framework

Cause: Installation of Oracle Grid Infrastructure fails when you run `root.sh`. Oracle Grid Infrastructure fails to start because the local host entry is missing from the hosts file.

The Oracle Grid Infrastructure `alert.log` file shows the following:

```
[/oracle/app/grid/bin/orarootagent.bin(11392)]CRS-5823:Could not initialize
agent framework. Details at (:CRSAGF00120:) in
/oracle/app/grid/log/node01/agent/crsd/orarootagent_root/orarootagent_root.log
2010-10-04 12:46:25.857
[ohasd(2401)]CRS-2765:Resource 'ora.crsd' has failed on server 'node01'.
```

You can verify this as the cause by checking `crsdOUT.log` file, and finding the following:

```
Unable to resolve address for localhost:2016
ONS runtime exiting
Fatal error: eONS: eonsapi.c: Aug 6 2009 02:53:02
```

Action: Add the local host entry in the hosts file.

Failed to connect to server, Connection refused by server, or Can't open display

Cause: These are typical of X Window display errors on Windows or UNIX systems, where `xhost` is not properly configured, or where you are running as a user account that is different from the account you used with the `startx` command to start the X server.

Action: In a local terminal window, log in as the user that started the X Window session, and enter the following command:

```
$ xhost fullyqualifiedRemoteHostname
```

For example:

```
$ xhost somehost.example.com
```

Then, enter the following commands, where *workstationname* is the host name or IP address of your workstation.

Bourne, Bash, or Korn shell:

```
$ DISPLAY=workstationname:0.0
$ export DISPLAY
```

To determine whether X Window applications display correctly on the local system, enter the following command:

```
$ xclock
```

The X clock should appear on your monitor. If `xclock` is not available, then install it on your system and repeat the test. If `xclock` is installed on your system, but the X clock fails to open on your display, then use of the `xhost` command may be restricted.

If you are using a VNC client to access the server, then ensure that you are accessing the visual that is assigned to the user that you are trying to use for the installation. For example, if you used the `su` command to become the installation owner on another user visual, and the `xhost` command use is restricted, then you cannot use the `xhost` command to change the display. If you use the visual assigned to the installation owner, then the correct display is available, and entering the `xclock` command results in the X clock starting on your display.

When the X clock appears, close the X clock, and start the installer again.

Failed to initialize ocrconfig

Cause: You have the wrong options configured for NFS in the `/etc/fstab` file.

You can confirm this by checking `ocrconfig.log` files located in the path `Grid_home/log/nodenummer/client` and finding the following:

```
/u02/app/crs/clusterregistry, ret -1, errno 75, os err string Value too large
for defined data type
2007-10-30 11:23:52.101: [ OCR OSD][3085960896]utopen:6': OCR location
```

Action: For file systems mounted on NFS, provide the correct mount configuration for NFS mounts in the `/etc/fstab` file:

```
rw, sync, bg, hard, nointr, tcp, vers=3, timeo=300, rsize=32768, wsize=32768, actimeo=0
```

Note: You should not have `netdev` in the mount instructions, or `vers=2`. The `netdev` option is only required for OCFS file systems, and `vers=2` forces the kernel to mount NFS using the earlier version 2 protocol.

After correcting the NFS mount information, remount the NFS mount point, and run the `root.sh` script again. For example, with the mount point `/u02`:

```
#umount /u02
#mount -a -t nfs
#cd $GRID_HOME
#sh root.sh
```

INS-32026 INSTALL_COMMON_HINT_DATABASE_LOCATION_ERROR

Cause: The location selected for the Grid home for a Cluster installation is located under an Oracle base directory.

Action: For Oracle Grid Infrastructure for a cluster installations, the Grid home must not be placed under one of the Oracle base directories, or under Oracle home directories of Oracle Database installation owners, or in the home directory of an installation owner. During installation, ownership of the path to the Grid home is changed to `root`. This change causes permission errors for other installations. In addition, the Oracle Clusterware software stack may not come up under an Oracle base path.

CLSRSC-444: Run root.sh command on the Node with OUI session

Cause: If this message appears listing a node that is not the one where you are running OUI, then the likely cause is that the named node shut down during or before the `root.sh` script completed its run.

Action: Complete running the `root.sh` script on all other cluster member nodes, and do not attempt to run the `root` script on the node named in the error message. After you complete Oracle Grid Infrastructure on all or part of the set of planned

cluster member nodes, start OUI and deinstall the failed Oracle Grid Infrastructure installation on the node named in the error. When you have deinstalled the failed installation on the node, add that node manually to the cluster.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to add a node

MEMORY_TARGET not supported on this system

Cause: On Linux systems, insufficient `/dev/shm` size for PGA and SGA.

If you are installing on a Linux system, note that Memory Size (SGA and PGA), which sets the initialization parameter `MEMORY_TARGET` or `MEMORY_MAX_TARGET`, cannot be greater than the shared memory file system (`/dev/shm`) on your operating system.

Action: Increase the `/dev/shm` mountpoint size. For example:

```
# mount -t tmpfs shmfs -o size=4g /dev/shm
```

Also, to make this change persistent across system restarts, add an entry in `/etc/fstab` similar to the following:

```
shmfs /dev/shm tmpfs size=7g 0 0
```

Nodes unavailable for selection from the OUI Node Selection screen

Cause: Oracle Grid Infrastructure is either not installed, or the Oracle Grid Infrastructure services are not up and running.

Action: Install Oracle Grid Infrastructure, or review the status of your installation. Consider restarting the nodes, as doing so may resolve the problem.

Node *nodename* is unreachable

Cause: Unavailable IP host

Action: Attempt the following:

1. Run the shell command `ifconfig -a`. Compare the output of this command with the contents of the `/etc/hosts` file to ensure that the node IP is listed.
2. Run the shell command `nslookup` to see if the host is reachable.

PROT-8: Failed to import data from specified file to the cluster registry

Cause: Insufficient space in an existing Oracle Cluster Registry device partition, which causes a migration failure while running `rootupgrade.sh`. To confirm, look for the error "utopen:12:Not enough space in the backing store" in the log file `$_GRID_HOME/log/hostname/client/ocrconfig_pid.log`, where `pid` stands for the process id.

Action: Identify a storage device that has 400 MB or more available space. Oracle recommends that you allocate the entire disk to Oracle ASM.

PRVE-0038 : The SSH LoginGraceTime setting, or fatal: Timeout before authentication

Cause: PRVE-0038: The SSH `LoginGraceTime` setting on node "*nodename*" may result in users being disconnected before login is completed. This error may be because the default timeout value for SSH connections is too low, or if the `LoginGraceTime` parameter is commented out.

Action: Oracle recommends uncommenting the `LoginGraceTime` parameter in the OpenSSH configuration file `/etc/ssh/sshd_config`, and setting it to a value of 0 (unlimited).

Timed out waiting for the CRS stack to start

Cause: If a configuration issue prevents the Oracle Grid Infrastructure software from installing successfully on all nodes, then you may see error messages such as "Timed out waiting for the CRS stack to start," or you may notice that Oracle Clusterware-managed resources were not create on some nodes after you exit the installer. You also may notice that resources have a status other than ONLINE.

Action: Unconfigure the Oracle Grid Infrastructure installation without removing binaries, and review log files to determine the cause of the configuration issue. After you have fixed the configuration issue, rerun the scripts used during installation to configure Oracle Clusterware.

See Also: [Section 10.5, "Unconfiguring Oracle Clusterware Without Removing Binaries"](#)

YPBINDPROC_DOMAIN: Domain not bound

Cause: This error can occur during postinstallation testing when the public network interconnect for a node is pulled out, and the VIP does not fail over. Instead, the node hangs, and users are unable to log in to the system. This error occurs when the Oracle home, `listener.ora`, Oracle log files, or any action scripts are located on an NAS device or NFS mount, and the name service cache daemon `nscd` has not been activated.

Action: Enter the following command on all nodes in the cluster to start the `nscd` service:

```
/sbin/service nscd start
```

A.2.1 Other Installation Issues and Errors

For additional help in resolving error messages, see My Oracle Support. For example, the note with Doc ID 1367631.1 contains some of the most common installation issues for Oracle Grid Infrastructure and Oracle Clusterware.

A.3 Interpreting CVU "Unknown" Output Messages Using Verbose Mode

If you run Cluster Verification Utility using the `-verbose` argument, and a Cluster Verification Utility command responds with `UNKNOWN` for a particular node, then this is because Cluster Verification Utility cannot determine if a check passed or failed. The following is a list of possible causes for an "Unknown" response:

- The node is down
- Common operating system command binaries required by Cluster Verification Utility are missing in the `/bin` directory in the Oracle Grid Infrastructure home or Oracle home directory
- The user account starting Cluster Verification Utility does not have privileges to run common operating system commands on the node
- The node is missing an operating system patch, or a required package
- The node has exceeded the maximum number of processes or maximum number of open files, or there is a problem with IPC segments, such as shared memory or semaphores

A.4 Interpreting CVU Messages About Oracle Grid Infrastructure Setup

If the Cluster Verification Utility (CVU) report indicates that your system fails to meet the requirements for Oracle Grid Infrastructure installation, then use the topics in this section to correct the problem or problems indicated in the report, and run CVU again.

- [User Equivalence Check Failed](#)
- [Node Reachability Check or Node Connectivity Check Failed](#)
- [User Existence Check or User-Group Relationship Check Failed](#)

User Equivalence Check Failed

Cause: Failure to establish user equivalency across all nodes. This can be due to not creating the required users, or failing to complete secure shell (SSH) configuration properly.

Action: Cluster Verification Utility provides a list of nodes on which user equivalence failed.

For each node listed as a failure node, review the installation owner user configuration to ensure that the user configuration is properly completed, and that SSH configuration is properly completed. The user that runs the Oracle Clusterware installation must have permissions to create SSH connections.

Oracle recommends that you use the SSH configuration option in OUI to configure SSH. You can use Cluster Verification Utility before installation if you configure SSH manually, or after installation, when SSH has been configured for installation.

For example, to check user equivalency for the user account `oracle`, use the command `su - oracle` and check user equivalence manually by running the `ssh` command on the local node with the `date` command argument using the following syntax:

```
$ ssh nodename date
```

The output from this command should be the timestamp of the remote node identified by the value that you use for `nodename`. If you are prompted for a password, then you need to configure SSH. If `ssh` is in the default location, the `/usr/bin` directory, then use `ssh` to configure user equivalence. You can also use `rsh` to confirm user equivalence.

If you see a message similar to the following when entering the `date` command with SSH, then this is the probable cause of the user equivalence error:

```
The authenticity of host 'node1 (140.87.152.153)' can't be established.  
RSA key fingerprint is 7z:ez:e7:f6:f4:f2:4f:8f:9z:79:85:62:20:90:92:z9.  
Are you sure you want to continue connecting (yes/no)?
```

Enter `yes`, and then run Cluster Verification Utility to determine if the user equivalency error is resolved.

If `ssh` is in a location other than the default, `/usr/bin`, then Cluster Verification Utility reports a user equivalence check failure. To avoid this error, navigate to the directory `Grid_home/cv/admin`, open the file `cvu_config` with a text editor, and add or update the key `ORACLE_SRVM_REMOTESHELL` to indicate the `ssh` path location on your system. For example:

```
# Locations for ssh and scp commands  
ORACLE_SRVM_REMOTESHELL=/usr/local/bin/ssh  
ORACLE_SRVM_REMOTECOPY=/usr/local/bin/scp
```

Note the following rules for modifying the `cvu_config` file:

- Key entries have the syntax `name=value`
- Each key entry and the value assigned to the key defines one property only
- Lines beginning with the number sign (#) are comment lines, and are ignored
- Lines that do not follow the syntax `name=value` are ignored

When you have changed the path configuration, run Cluster Verification Utility again. If `ssh` is in another location than the default, you also need to start OUI with additional arguments to specify a different location for the remote shell and remote copy commands. Enter `runInstaller -help` to obtain information about how to use these arguments.

Note: When you or OUI run `ssh` or `rsh` commands, including any login or other shell scripts they start, you may see errors about invalid arguments or standard input if the scripts generate any output. You should correct the cause of these errors.

To stop the errors, remove all commands from the `oracle` user's login scripts that generate output when you run `ssh` or `rsh` commands.

If you see messages about X11 forwarding, then complete the task [Section 6.2.4, "Setting Remote Display and X11 Forwarding Configuration"](#) to resolve this issue.

If you see errors similar to the following:

```
stty: standard input: Invalid argument
stty: standard input: Invalid argument
```

These errors are produced if hidden files on the system (for example, `.bashrc` or `.cshrc`) contain `stty` commands. If you see these errors, then see [Section 6.2.5, "Preventing Installation Errors Caused by Terminal Output Commands"](#) to correct the cause of these errors.

Node Reachability Check or Node Connectivity Check Failed

Cause: One or more nodes in the cluster cannot be reached using TCP/IP protocol, through either the public or private interconnects.

Action: Use the command `/bin/ping address` to check each node address. When you find an address that cannot be reached, check your list of public and private addresses to make sure that you have them correctly configured. If you use third-party vendor clusterware, then see the vendor documentation for assistance. Ensure that the public and private network interfaces have the same interface names on each node of your cluster.

User Existence Check or User-Group Relationship Check Failed

Cause: The administrative privileges for users and groups required for installation are missing or incorrect.

Action: Use the `id` command on each node to confirm that the installation owner user (for example, `grid` or `oracle`) is created with the correct group membership. Ensure that you have created the required groups, and create or modify the user account on affected nodes to establish required group membership.

See Also: [Section 6.1, "Creating Groups, Users and Paths for Oracle Grid Infrastructure"](#) for instructions about how to create required groups, and how to configure the installation owner user

A.5 About the Oracle Clusterware Alert Log

Oracle Clusterware uses Oracle Database fault diagnosability infrastructure to manage diagnostic data and its alert log. As a result, most diagnostic data resides in the Automatic Diagnostic Repository (ADR), a collection of directories and files located under a base directory that you specify during installation. Starting with Oracle Clusterware 12c release 1 (12.1.0.2), diagnostic data files written by Oracle Clusterware programs are known as trace files and have a .trc file extension, and appear together in the trace subdirectory of the ADR home. Besides trace files, the trace subdirectory in the Oracle Clusterware ADR home contains the simple text Oracle Clusterware alert log. It always has the name alert.log. The alert log is also written as an XML file in the alert subdirectory of the ADR home, but the text alert log is most easily read.

The Oracle Clusterware alert log is the first place to look for serious errors. In the event of an error, it can contain path information to diagnostic logs that can provide specific information about the cause of errors.

After installation, Oracle Clusterware posts alert messages when important events occur. For example, you may see alert messages from the Cluster Ready Services daemon process (CRSD) when it starts, if it aborts, if the failover process fails, or if automatic restart of an Oracle Clusterware resource fails.

Oracle Enterprise Manager monitors the Oracle Clusterware log file and posts an alert on the Cluster Home page if an error is detected. For example, if a voting file is not available, a CRS-1604 error is raised, and a critical alert is posted on the Cluster Home page. You can customize the error detection and alert settings on the Metric and Policy Settings page.

The location of the Oracle Clusterware log file is `ORACLE_BASE/diag/crs/hostname/crs/trace/alert.log`, where `ORACLE_BASE` is the Oracle base path you specified when you installed Oracle Grid Infrastructure and `hostname` is the name of the host.

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware troubleshooting
- *Oracle Database Utilities Guide* for information about the Automatic Diagnostic Repository Command Interpreter (ADRCI) utility to manage Oracle Database diagnostic data
- *Oracle Database Administrator's Guide* for more information about managing diagnostic data

A.6 Missing Operating System Packages On Linux

You have missing operating system packages on your system if you receive error messages such as the following during Oracle Grid Infrastructure, Oracle RAC, or Oracle Database installation:

```
libstdc++.so.5: cannot open shared object file: No such file or directory
libXp.so.6: cannot open shared object file: No such file or directory
```

Errors such as these should not occur, as missing packages should have been identified during installation. They may indicate that you are using an operating system distribution that has not been certified, or that you are using an earlier version of the Cluster Verification Utility.

If you have a Linux support network configured, such as the Red Hat network or Oracle Unbreakable Linux support, then use the `up2date` command to determine the name of the package. For example:

```
# up2date --whatprovides libstdc++.so.5
compat-libstdc++-33.3.2.3-47.3
```

Also, download the most recent version of Cluster Verification Utility to make sure that you have the current required packages list. You can obtain the most recent version at the following URL:

<http://www.oracle.com/technetwork/database/options/clustering/downloads/cvu-download-homepage-099973.html>

A.7 Performing Cluster Diagnostics During Oracle Grid Infrastructure Installations

If the installer does not display the Node Selection page, then use the following command syntax to check the integrity of the Cluster Manager:

```
cluvfy comp clumgr -n node_list -verbose
```

In the preceding syntax example, the variable `node_list` is the list of nodes in your cluster, separated by commas.

Note: If you encounter unexplained installation errors during or after a period when cron jobs are run, then your cron job may have deleted temporary files 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.

A.8 About Using CVU Cluster Healthchecks After Installation

Starting with Oracle Grid Infrastructure 11g Release 2 (11.2.0.3) and later, you can use the Cluster Verification Utility (CVU) healthcheck command option to check your Oracle Clusterware and Oracle Database installations for their compliance with mandatory requirements and best practices guidelines, and to check to ensure that they are functioning properly.

Use the following syntax to run the healthcheck command option:

```
cluvfy comp healthcheck [-collect {cluster|database}] [-db db_unique_name]
[-bestpractice|-mandatory] [-deviations] [-html] [-save [-savedir directory_path]]
```

For example:

```
$ cd /home/grid/cvu_home/bin
$ ./cluvfy comp healthcheck -collect cluster -bestpractice -deviations -html
```

The options are:

- `-collect [cluster|database]`

Use this flag to specify that you want to perform checks for Oracle Clusterware (-cluster) or Oracle Database (-database). If you do not use the -collect flag with the healthcheck option, then `cluvfy comp healthcheck` performs checks for both Oracle Clusterware and Oracle Database.

- `-db db_unique_name`

Use this flag to specify checks on the database unique name that you enter after the -db flag.

CVU uses JDBC to connect to the database as the user `cvusys` to verify various database parameters. For this reason, if you want checks to be performed for the database you specify with the -db flag, then you must first create the `cvusys` user on that database, and grant that user the CVU-specific role, `cvusapp`. You must also grant members of the `cvusapp` role `select` permissions on system tables.

A SQL script is included in `CVU_home/cv/admin/cvusys.sql` to facilitate the creation of this user. Use this SQL script to create the `cvusys` user on all the databases that you want to verify using CVU.

If you use the -db flag but do not provide a database unique name, then CVU discovers all the Oracle Databases on the cluster. If you want to perform best practices checks on these databases, then you must create the `cvusys` user on each database, and grant that user the `cvusapp` role with the `select` privileges needed to perform the best practice checks.

- `[-bestpractice | -mandatory] [-deviations]`

Use the -bestpractice flag to specify best practice checks, and the mandatory flag to specify mandatory checks. Add the -deviations flag to specify that you want to see only the deviations from either the best practice recommendations or the mandatory requirements. You can specify either the -bestpractice or -mandatory flag, but not both flags. If you specify neither -bestpractice or -mandatory, then both best practices and mandatory requirements are displayed.

- `-html`

Use the -html flag to generate a detailed report in HTML format.

If you specify the -html flag, and a browser CVU recognizes is available on the system, then the browser is started and the report is displayed on the browser when the checks are complete.

If you do not specify the html flag, then the detailed report is generated in a text file.

- `-save [-savedir dir_path]`

Use the -save or -save -savedir flags to save validation reports (`cvuchekdkreport_timestamp.txt` and `cvucheckreport_timestamp.htm`), where `timestamp` is the time and date of the validation report.

If you use the -save flag by itself, then the reports are saved in the path `CVU_home/cv/report`, where `CVU_home` is the location of the CVU binaries.

If you use the flags -save -savedir, and enter a path where you want the CVU reports saved, then the CVU reports are saved in the path you specify.

A.9 Interconnect Configuration Issues

If you plan to use multiple network interface cards (NICs) for the interconnect, and you do not configure them during installation or after installation with Redundant

Interconnect Usage, then you should use a third-party solution to bond the interfaces at the operating system level. Otherwise, the failure of a single NIC will affect the availability of the cluster node.

If you install Oracle Grid Infrastructure and Oracle RAC, then they must use the same NIC or bonded NIC cards for the interconnect.

If you use bonded NIC cards, and use the Oracle Clusterware Redundant Interconnect Usage feature, then they should be on different subnets. If you use a third-party vendor method of aggregation, such as bonding or IPMP, then follow the directions for that vendor's product.

If you encounter errors, then carry out the following system checks:

- Verify with your network providers that they are using correct cables (length, type) and software on their switches. In some cases, to avoid bugs that cause disconnects under loads, or to support additional features such as Jumbo Frames, you may need a firmware upgrade on interconnect switches, or you may need newer NIC driver or firmware at the operating system level. Running without such fixes can cause later instabilities to Oracle RAC databases, even though the initial installation seems to work.
- Review VLAN configurations, duplex settings, and auto-negotiation in accordance with vendor and Oracle recommendations.

A.10 SCAN VIP and SCAN Listener Issues

If the final check of your installation reports errors related to the SCAN VIP addresses or listeners, then check the following items to make sure your network is configured correctly:

- Check the file `/etc/resolv.conf` - verify the contents are the same on each node
- Verify that there is a DNS entry for the SCAN, and that it resolves to three valid IP addresses. Use the command `nslookup scan-name`; this command should return the DNS server name and the three IP addresses configured for the SCAN.
- Use the `ping` command to test the IP addresses assigned to the SCAN; you should receive a response for each IP address.

Note: If you do not have a DNS configured for your cluster environment, then you can create an entry for the SCAN in the `/etc/hosts` file on each node. However, using the `/etc/hosts` file to resolve the SCAN results in having only one SCAN available for the entire cluster instead of three. Only the first entry for SCAN in the `hosts` file is used.

- Ensure the SCAN VIP uses the same netmask that is used by the public interface.

If you need additional assistance troubleshooting errors related to the SCAN, SCAN VIP or listeners, then refer to My Oracle Support. For example, the note with Doc ID 1373350.1 contains some of the most common issues for the SCAN VIPs and listeners.

A.11 Storage Configuration Issues

The following is a list of issues involving storage configuration:

- [Recovery from Losing a Node Filesystem or Grid Home](#)

- [Oracle ASM Library Driver Issues](#)
- [Oracle ASM Issues After Upgrading Oracle Grid Infrastructure](#)
- [Oracle ASM Issues After Downgrading Oracle Grid Infrastructure for Standalone Server \(Oracle Restart\)](#)

A.11.1 Recovery from Losing a Node Filesystem or Grid Home

If you remove a filesystem by mistake, or encounter another storage configuration issue that results in losing the Oracle Local Registry or otherwise corrupting a node, you can recover the node in one of two ways:

- Restore the node from an operating system level backup
- Remove the node, and then add the node, using *Grid home*/addnode/addnode.sh. Profile information for is copied to the node, and the node is restored.

Using addnode.sh enables cluster nodes to be removed and added again, so that they can be restored from the remaining nodes in the cluster. If you add nodes in a GNS configuration, then that is called Grid Plug and Play (GPnP). GPnP uses profiles to configure nodes, which eliminates configuration data requirements for nodes and the need for explicit add and delete nodes steps. GPnP allows a system administrator to take a template system image and run it on a new node with no further configuration. GPnP removes many manual operations, reduces the opportunity for errors, and encourages configurations that can be changed easily. Removal of individual node configuration makes the nodes easier to replace, because nodes do not need to contain individually-managed states.

GPnP reduces the cost of installing, configuring, and managing database nodes by making their node state disposable. It allows nodes to be easily replaced with a regenerated state.

See Also: *Oracle Clusterware Administration and Deployment Guide* for information about how to add nodes manually or with GNS

A.11.2 Oracle ASM Library Driver Issues

The following is a list of Oracle ASM driver library error messages, and how to address these errors:

Asmtool: Unable to clear device "devicepath": Input/output error

Cause: This is a write access error that can have several causes.

Action: If the disk is mounted, then unmount it. For example: `umount /dev/sdb1`. Ensure that the group and user that owns the device is the Oracle Grid Infrastructure installation owner and the oraInventory group. For example: `chown grid:oinstall`.

Unable to open ASMLib; Unable to find candidate disks 'ORCL:*

Cause: If you have created disks but you cannot discover candidate disks from OUI, this could be due to a variety of configuration errors that prevent access to Oracle ASM storage.

Action: As the Oracle Grid Infrastructure installation owner, enter the command `/usr/sbin/oracleasm-discover`, using the ASM disk path *asm_diskstring*. For example:

```
[grid@node1]$ /usr/sbin/oracleasm-discover 'ORCL:*
```

If you do not have `/usr/sbin/oracleasm-discover`, then you do not have `oracleasm` installed. If you do have the command, then you should be able to determine if ASMLib is enabled, if disks are created, and if other tasks to create candidate disks are completed.

If you have resolved the issue, then you should see output similar to the following when you enter the command:

```
[grid@node1]$ /usr/sbin/oracleasm-discover 'ORCL:*'

Using ASMLib from /opt/oracle/extapi/64/asm/orcl/1/libasm.so

[ASM Library - Generic Linux, version 2.0.4 (KABI_V2)]
Discovered disk: ORCL:DISK1 [78140097 blocks (40007729664 bytes), maxio 512]
Discovered disk: ORCL:DISK2 [78140097 blocks (40007729664 bytes), maxio 512]
Discovered disk: ORCL:DISK3 [78140097 blocks (40007729664 bytes), maxio 512]
```

A.11.3 Oracle ASM Issues After Upgrading Oracle Grid Infrastructure

The following section explains an error that can occur when you upgrade Oracle Grid Infrastructure, and how to address it:

CRS-0219: Could not update resource 'ora.node1.asm1.inst'

Cause: After upgrading Oracle Grid Infrastructure, Oracle ASM client databases prior to Oracle Database 12c are unable to obtain the Oracle ASM instance aliases on the `ora.asm` resource through the `ALIAS_NAME` attribute.

Action: You must use Local ASM or set the cardinality for Flex ASM to ALL, instead of the default of 3. Use the following command to modify the Oracle ASM resource (`ora.asm`):

```
$ srvctl modify asm -count ALL
```

This setting changes the cardinality so that Flex ASM instances run on all nodes.

See Also: [Section 9.3.3, "Making Oracle ASM Available to Earlier Oracle Database Releases"](#) for information about making Oracle ASM available to Oracle Database releases earlier than 12c Release 1

A.11.4 Oracle ASM Issues After Downgrading Oracle Grid Infrastructure for Standalone Server (Oracle Restart)

The following section explains an error that can occur when you downgrade Oracle Grid Infrastructure for standalone server (Oracle Restart), and how to address it:

CRS-2529: Unable to act on 'ora.cssd' because that would require stopping or relocating 'ora.asm'

Cause: After downgrading Oracle Grid Infrastructure for a standalone server (Oracle Restart) from 12.1.0.2 to 12.1.0.1, the `ora.asm` resource does not contain the Server Parameter File (SPFILE) parameter.

Action: When you downgrade Oracle Grid Infrastructure for a standalone server (Oracle Restart) from 12.1.0.2 to 12.1.0.1, you must explicitly add the Server Parameter File (SPFILE) from the `ora.asm` resource when adding the Oracle ASM resource for 12.1.0.1.

Follow these steps when you downgrade Oracle Restart from 12.1.0.2 to 12.1.0.1:

1. In your 12.1.0.2 Oracle Restart installed configuration, query the `SPFILE` parameter from the Oracle ASM resource (`ora.asm`) and remember it:

```
srvctl config asm
```

2. Deconfigure the 12.1.0.2 release Oracle Restart:

```
Grid_home/crs/install/roothas.pl -deconfig -force
```

3. Install the 12.1.0.1 release Oracle Restart by running `root.sh`:

```
$ Grid_home/root.sh
```

4. Add the listener resource:

```
$ Grid_home/bin/srvctl add LISTENER
```

5. Add the Oracle ASM resource and provide the `SPFILE` parameter for the 12.1.0.2 Oracle Restart configuration obtained in step 1:

```
$ Grid_home/bin/srvctl add asm  
[-spfile <spfile>] [-diskstring <asm_diskstring>]
```

See Also: *Oracle Database Installation Guide* for information about installing and deconfiguring Oracle Restart

A.12 Failed or Incomplete Installations and Upgrades

During installations or upgrades of Oracle Grid Infrastructure, the following actions take place:

1. Oracle Universal Installer (OUI) accepts inputs to configure Oracle Grid Infrastructure software on your system.
2. You are instructed to run either the `oraInstRoot.sh` or `root.sh` script or both.
3. You run the scripts either manually or through root automation.
4. OUI runs configuration assistants. The Oracle Grid Infrastructure software installation completes successfully.

If OUI exits before the `root.sh` or `rootupgrade.sh` script runs, or if OUI exits before the installation or upgrade session is completed successfully, then the Oracle Grid Infrastructure installation or upgrade is incomplete. If your installation or upgrade does not complete, then Oracle Clusterware does not work correctly. If you are performing an upgrade, then an incomplete upgrade can result in some nodes being upgraded to the latest software and others nodes not upgraded at all. If you are performing an installation, the incomplete installation can result in some nodes not being a part of the cluster.

Additionally, from Oracle Grid Infrastructure release 11.2.0.3 or later, the following messages may be seen during installation or upgrade:

```
ACFS-9427 Failed to unload ADVM/ACFS drivers. A system reboot is  
recommended
```

```
ACFS-9428 Failed to load ADVM/ACFS drivers. A system reboot is recommended
```

```
CLSRSC-400: A system reboot is required to continue installing
```

To resolve this error, you must reboot the server, and then follow the steps for completing an incomplete installation or upgrade as documented in the following sections:

- [Completing Failed or Interrupted Upgrades](#)
- [Completing Failed or Interrupted Installations](#)

A.12.1 Completing Failed or Interrupted Upgrades

If 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, the upgrade remains incomplete. In an incomplete 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 upgrade on the affected nodes manually.

This section contains the following tasks:

- [Continuing Upgrade When Force Upgrade in Rolling Upgrade Mode Fails](#)
- [Continuing Upgrade When Upgrade Fails on the First Node](#)
- [Continuing Upgrade When Upgrade Fails on Nodes Other Than the First Node](#)

A.12.1.1 Continuing Upgrade When Force Upgrade in Rolling Upgrade Mode Fails

If you attempt to force upgrade cluster nodes in the rolling upgrade mode, you may see the following error:

CRS 1137 - Rejecting the rolling upgrade mode change because the cluster was forcibly upgraded.

Cause: The rolling upgrade mode change was rejected because the cluster was forcibly upgraded.

Action: Delete the nodes that were not upgraded using the procedure documented in *Oracle Clusterware Administration and Deployment Guide*. You can then retry the rolling upgrade process using the `crsctl start rollingupgrade` command as documented in [Section B.8, "Performing Rolling Upgrade of Oracle Grid Infrastructure"](#).

A.12.1.2 Continuing Upgrade When Upgrade Fails on the First Node

When the first node cannot be upgraded, do the following:

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. Run the `rootupgrade.sh` script on that node again.
2. Complete the upgrade of all other nodes in the cluster.
3. Configure a response file, and provide passwords for the installation. See [Section C.5, "Postinstallation Configuration Using a Response File"](#) for information about how to create the response file.
4. To complete the upgrade, log in as the Grid installation owner, and run the script `configToolAllCommands`, located in the path `Gridhome/cfgtoollogs/configToolAllCommands`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ cd /u01/app/12.1.0/grid/cfgtoollogs
[grid@node1]$ ./configToolAllCommands RESPONSE_FILE=gridinstall.rsp
```

A.12.1.3 Continuing Upgrade When Upgrade Fails on Nodes Other Than the First Node

For nodes other than the first node (the node on which the upgrade was started):

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/12.1.0/grid
[root@node6]# ./rootupgrade.sh
```

A.12.2 Completing Failed or Interrupted Installations

If OUI exits on the node from which you started the install, or the node reboots before you confirm that the `orainstRoot.sh` or `root.sh` script were run on all nodes, the install remains incomplete. In an incomplete install, 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 install on the affected nodes manually.

This section contains the following tasks:

- [Continuing Incomplete Installations on First Node](#)
- [Continuing Installation on Nodes Other Than the First Node](#)

A.12.2.1 Continuing Incomplete Installations on First Node

The first node must finish installation before the rest of the clustered nodes. To continue an incomplete installation 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. 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. See [Section C.5, "Postinstallation Configuration Using a Response File"](#) for information about how to create the response file.

6. To complete the installation, log in as the Grid installation owner, and run the script `configToolAllCommands`, located in the path `Gridhome/cfgtoollogs/configToolAllCommands`, specifying the response file that you created. For example, where the response file is `gridinstall.rsp`:

```
[grid@node1]$ cd /u01/app/12.1.0/grid/cfgtoollogs
[grid@node1]$ ./configToolAllCommands RESPONSE_FILE=gridinstall.rsp
```

A.12.2.2 Continuing Installation on Nodes Other Than the First Node

For nodes other than the first node (the node on which the installation was started):

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/12.1.0/grid
[root@node6]# ./root.sh
```

4. Continue the installation from the OUI instance on the first node.

How to Upgrade to Oracle Grid Infrastructure 12c Release 1

This appendix describes how to perform Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) upgrades.

Oracle Clusterware upgrades can be rolling upgrades, in which a subset of nodes are brought down and upgraded while other nodes remain active. Oracle ASM 12c Release 1 (12.1) upgrades can be rolling upgrades. If you upgrade a subset of nodes, then a software-only installation is performed on the existing cluster nodes that you do not select for upgrade.

This appendix contains the following topics:

- [Back Up the Oracle Software Before Upgrades](#)
- [About Oracle Grid Infrastructure and Oracle ASM Upgrade and Downgrade](#)
- [Options for Oracle Grid Infrastructure Upgrades and Downgrades](#)
- [Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades](#)
- [Preparing to Upgrade an Existing Oracle Clusterware Installation](#)
- [Using CVU to Validate Readiness for Oracle Clusterware Upgrades](#)
- [Understanding Rolling Upgrades Using Batches](#)
- [Performing Rolling Upgrade of Oracle Grid Infrastructure](#)
- [Restrictions and Guidelines for Upgrading and Patching Oracle ASM](#)
- [Performing Rolling Upgrade of Oracle ASM](#)
- [Applying Patches to Oracle ASM](#)
- [Updating Oracle Enterprise Manager Cloud Control Target Parameters](#)
- [Unlocking the Existing Oracle Clusterware Installation](#)
- [Checking Cluster Health Monitor Repository Size After Upgrading](#)
- [Downgrading Oracle Clusterware After an Upgrade](#)

B.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.

B.2 About Oracle Grid Infrastructure and Oracle ASM Upgrade and Downgrade

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 1 (12.1) to prior releases of Oracle Grid Infrastructure. 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.

If you have an existing Oracle ASM 11g Release 1 (11.1) or 10g release instance, with Oracle ASM in a separate home, then you can either upgrade it at the time that you install Oracle Grid Infrastructure, or you can upgrade it after the installation, using Oracle ASM Configuration Assistant (ASMCA). However, be aware that a number of Oracle ASM features are disabled until you upgrade Oracle ASM, and Oracle Clusterware management of Oracle ASM does not function correctly until Oracle ASM is upgraded, because Oracle Clusterware only manages Oracle ASM when it is running in the Oracle Grid Infrastructure home. For this reason, Oracle recommends that if you do not upgrade Oracle ASM at the same time as you upgrade Oracle Clusterware, then you should upgrade Oracle ASM immediately afterward. This issue does not apply to Oracle ASM 11g Release 2 (11.2) and later, as the Oracle Grid Infrastructure home contains Oracle ASM binaries as well.

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

B.3 Options for Oracle Grid Infrastructure Upgrades and Downgrades

Upgrade options from Oracle Grid Infrastructure 11g to Oracle Grid Infrastructure 12c 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

Upgrade options from Oracle Grid Infrastructure 11g Release 2 (11.2) to Oracle Grid Infrastructure 12c include the following:

- Oracle Grid Infrastructure rolling upgrade, with OCR and voting disks on Oracle ASM
- Oracle Grid Infrastructure complete cluster upgrade (downtime, non-rolling), with OCR and voting disks on Oracle ASM

Upgrade options from releases before Oracle Grid Infrastructure 11g Release 2 (11.2) to Oracle Grid Infrastructure 12c include the following:

- Oracle Grid Infrastructure rolling upgrade, with OCR and voting disks on storage other than Oracle ASM or shared file system
- Oracle Grid Infrastructure complete cluster upgrade (downtime, non-rolling), with OCR and voting disks on storage other than Oracle ASM or shared file system

Downgrade options from Oracle Grid Infrastructure 12c to earlier releases include the following:

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2)
- Oracle Grid Infrastructure downgrades to releases before Oracle Grid Infrastructure 11g Release 2 (11.2), Oracle Grid Infrastructure 11g Release 1 (11.1), Oracle Clusterware and Oracle ASM 10g, if storage for OCR and voting files is on storage other than Oracle ASM

B.4 Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades

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 11g Release 2 (11.2) or Oracle RAC database 11g Release 2 (11.2) installations.

See Also: [Section B.6, "Using CVU to Validate Readiness for Oracle Clusterware Upgrades"](#)

Be aware of the following restrictions and changes for upgrades to Oracle Grid Infrastructure installations, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM):

- When you upgrade from Oracle Grid Infrastructure 11g or Oracle Clusterware and Oracle ASM 10g releases to Oracle Grid Infrastructure 12c Release 1 (12.1), you upgrade to a standard cluster configuration. You can enable Oracle Flex Cluster configuration after the upgrade.
- If the Oracle Cluster Registry (OCR) and voting file locations for your current installation are on raw or block devices, then you must migrate them to Oracle ASM disk groups or shared file systems before upgrading to Oracle Grid Infrastructure 12c.
- 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, and you want to migrate these files to Oracle ASM rather than to a shared file system, then you must upgrade to Oracle Grid Infrastructure 11g Release 2 (11.2) before you upgrade to Oracle Grid Infrastructure 12c.

- Downgrades from an Oracle Grid Infrastructure 12c Release 1 (12.1) Oracle Flex Cluster configuration to a Standard cluster configuration are not supported. All cluster configurations in releases earlier than Oracle Grid Infrastructure 12c are Standard cluster configurations. This downgrade restriction includes downgrades from an Oracle Flex Cluster to Oracle Grid Infrastructure 11g cluster, or to Oracle Clusterware and Oracle ASM 10g clusters.
- You can 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 (12.1), you can only downgrade to Oracle Grid Infrastructure 11g Release 2 (11.2).
- 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 a standard configuration Oracle Grid Infrastructure 12c cluster, your release must be greater than or equal to Oracle Clusterware 10g Release 1 (10.1.0.5), Oracle Clusterware 10g Release 2 (10.2.0.3), Oracle Grid Infrastructure 11g Release 1 (11.1.0.6), or Oracle Grid Infrastructure 11g Release 2 (11.2).
- To upgrade existing Oracle Grid Infrastructure installations from Oracle Grid Infrastructure 11g Release 2 (11.2.0.2) to a later release, you must apply patch 11.2.0.2.3 (11.2.0.2 PSU 3) or later.
- 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 1 (12.1), you must first verify if you need to apply any mandatory patches for upgrade to succeed. See [Section B.6](#) for steps to check readiness.

See Also: Oracle 12c Upgrade Companion (My Oracle Support Note 1462240.1):

<https://support.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1462240.1>

- Oracle Clusterware and Oracle ASM upgrades are always out-of-place upgrades. You cannot perform an in-place upgrade of Oracle Clusterware and Oracle ASM to existing homes.
- If the existing Oracle Clusterware home is a shared home, note that you can use a non-shared home for the Oracle Grid Infrastructure for a cluster home for Oracle Clusterware and Oracle ASM 12c Release 1 (12.1).
- The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 12c Release 1 (12.1) upgrade. Before Oracle Database 11g, either all Oracle software installations were owned by the Oracle user, typically `oracle`, or Oracle Database software was owned by `oracle`, and Oracle Clusterware software was owned by a separate user, typically `crs`.

- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.
- During a major release upgrade to Oracle Grid Infrastructure 12c Release 1 (12.1), the software in the 12c Release 1 (12.1) 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.
- You can perform upgrades on a shared Oracle Clusterware home.
- During Oracle Clusterware installation, if there is a single instance Oracle ASM release on the local node, then it is converted to a clustered Oracle ASM 12c Release 1 (12.1) 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 is being performed), then it will remain a single instance Oracle ASM installation. However, during installation, if you select to place the Oracle Cluster Registry (OCR) and voting files on Oracle ASM, then a clustered Oracle ASM installation is created on all nodes in the cluster, and the single instance Oracle ASM installation on the remote node will become nonfunctional.
- After completing the force upgrade of a cluster to a release, all inaccessible nodes must be deleted from the cluster or joined to the cluster before starting the cluster upgrade to a later release.

See Also: *Oracle Database Upgrade Guide* for additional information about preparing for upgrades

B.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 sections list the steps you can perform before you upgrade Oracle Grid Infrastructure:

- [Checks to Complete Before Upgrading Oracle Clusterware](#)
- [Unset Oracle Environment Variables](#)
- [Running the Oracle ORAchk Upgrade Readiness Assessment](#)

B.5.1 Checks to Complete Before Upgrading Oracle Clusterware

Complete the following tasks before starting an upgrade:

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 will help you to ensure all required prerequisites are met.

Ensure that you have information you will 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, as described in [Chapter 5](#).
 - Privileged user operating system groups, as described in [Chapter 6](#).
 - root user access, to run scripts as root during installation, using one of the options described in [Section 8.1.1](#).
2. 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:

```
$ unset ORACLE_BASE
$ unset ORACLE_HOME
$ unset ORACLE_SID
```

See Also: [Section B.5.2, "Unset Oracle Environment Variables"](#)

3. If the cluster was previously forcibly upgraded, then ensure that all inaccessible nodes have been deleted from the cluster or joined to the cluster before starting another upgrade. For example, if the cluster was forcibly upgraded from 11.2.0.3 to 12.1.0.1, then ensure that all inaccessible nodes have been deleted from the cluster or joined to the cluster before upgrading to another release, for example, 12.1.0.2.

B.5.2 Unset Oracle Environment Variables

Unset Oracle environment variables.

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 installation owner login shell profiles (for example, .profile or .cshrc) do not have ORA_CRS_HOME set.

If you have had 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.

Also, ensure that the \$ORACLE_HOME/bin path is removed from your PATH environment variable.

B.5.3 Running the Oracle ORAchk Upgrade Readiness Assessment

ORAchk (Oracle RAC Configuration Audit Tool) 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, and 12.1.0.2. 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 configuration audit tool, refer to My Oracle Support note 1457357.1, which is available at the following URL:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1457357.1>

B.6 Using CVU to Validate Readiness for Oracle Clusterware Upgrades

You can use Cluster Verification Utility (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.

This section contains the following topics:

- [About the CVU Grid Upgrade Validation Command Options](#)
- [Example of Verifying System Upgrade Readiness for Grid Infrastructure](#)

B.6.1 About the CVU Grid Upgrade Validation Command Options

You can run upgrade validations in one of two ways:

- Run OUI, and allow the CVU validation built into OUI to perform system checks and generate fixup scripts
- Run the CVU manual script `cluvfy.sh` script 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 staging area for the upgrade, where the `runcluvfy.sh` command is located, 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][-method {sudo|root} [-location dir_path] [-user user_name]] [-verbose]
```

The options are:

- `-rolling`
Use this flag to verify readiness for rolling upgrades.
- `-src_crshome src_Gridhome`
Use this flag 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 flag 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` flag 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: 12.1.0.1.0.

- `-fixup [-method {sudo|root}] [-location dir_path] [-user user_name]`

Use the `-fixup` flag 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` flag 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` flag to provide the path to Sudo on the server, and enter the `-user` flag to provide the user account with Sudo privileges.

- `-verbose`

Use the `-verbose` flag to produce detailed output of individual checks.

B.6.2 Example of Verifying System Upgrade Readiness for Grid Infrastructure

You can verify that the permissions required for installing Oracle Clusterware have been configured 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/12.1.0/grid -dest_version
12.1.0.1 -fixup -verbose
```

See Also: *Oracle Database Upgrade Guide*

B.7 Understanding Rolling Upgrades Using Batches

Upgrades from earlier releases require that 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. If you upgrade from Oracle Grid Infrastructure 11g Release 11.2.0.2 and later to any later release of Oracle Grid Infrastructure, then all nodes are selected for upgrade by default.

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.

B.8 Performing Rolling Upgrade of Oracle Grid Infrastructure

This section contains the following topics:

- [Performing a Standard Upgrade from an Earlier Release](#)
- [Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable](#)
- [Upgrading Inaccessible Nodes After Forcing an Upgrade](#)

B.8.1 Performing a Standard Upgrade from an Earlier Release

Use the following procedure to upgrade the cluster from an earlier release:

1. Start the installer, and select the option to upgrade an existing Oracle Clusterware and Oracle ASM installation.
2. On the node selection page, select all nodes.
3. 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.
4. 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 12.1.0.1 Oracle Flex Cluster, Oracle recommends that you run the `rootupgrade.sh` script on all Hub Nodes before running it on Leaf Nodes.

5. After running the `rootupgrade.sh` script on the last node in the cluster, if you are upgrading from a release earlier than Oracle Grid Infrastructure 11g Release 2 (11.2.0.2), and left the check box labeled **ASMCA** checked, as is the default, then Oracle Automatic Storage Management Configuration Assistant ASMCA runs automatically, and the Oracle Grid Infrastructure upgrade is complete. If you unchecked the box during the interview stage of the upgrade, then ASMCA is not run automatically.

If an earlier release of Oracle Automatic Storage Management (Oracle ASM) is installed, then the installer starts ASMCA to upgrade Oracle ASM to 12c Release 1 (12.1). You can choose to upgrade Oracle ASM at this time, or upgrade it later.

Oracle recommends that you upgrade Oracle ASM at the same time that you upgrade Oracle Clusterware. Until Oracle ASM is upgraded, Oracle Databases

that use Oracle ASM cannot be created and the Oracle ASM management tools in the Oracle Grid Infrastructure 12c Release 1 (12.1) home (for example, `srvctl`) do not work.

6. 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.

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.

See Also: [Section A.12, "Failed or Incomplete Installations and Upgrades"](#) for information about completing failed or incomplete upgrades

B.8.2 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` command with the `-force` flag on any of the nodes where the `rootupgrade.sh` script has already completed as follows:

```
Grid_home/rootupgrade.sh -force
```

For example:

```
# /u01/app/12.1.0/grid/rootupgrade.sh -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.
- If the cluster was previously forcibly upgraded, then ensure that all inaccessible nodes have been deleted from the cluster or joined to the cluster before starting the upgrade.

B.8.3 Upgrading Inaccessible Nodes After Forcing an Upgrade

Starting with Oracle Grid Infrastructure 12c, after you complete a force cluster upgrade, you can join inaccessible nodes to the cluster as an alternative to deleting the nodes, which was required in earlier releases. To use this option, Oracle Grid Infrastructure 12c Release 1 (12.1) software must already be installed on the nodes.

To complete the upgrade of nodes that were inaccessible or unreachable:

1. Log in as the Grid user on the node you want to join to the cluster.
2. Change directory to the `/crs/install` directory in the Oracle Grid Infrastructure 12c Release 1 (12.1) Grid home. For example:

```
$ cd /u01/12.1.0/grid/crs/install
```

3. Run the following PERL command, where *existingnode* is the name of the option and *upgraded_node* is any node that was successfully upgraded and is currently part of the cluster:

```
$ rootupgrade.sh -join -existingnode upgraded_node
```

Note: The `-join` operation is not supported for Oracle Clusterware releases earlier than 11.2.0.1.0. In such cases, delete the node and add it to Oracle Clusterware using the `addNode` command.

B.8.4 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, run the following command:

```
rootupgrade.sh -force -first
```

B.9 Restrictions and Guidelines for Upgrading and Patching Oracle ASM

Note the following if you intend to perform either full release or software patch level rolling upgrades of Oracle ASM:

- The active release of Oracle Clusterware must be 12c Release 1 (12.1). To determine the active release, enter the following command:


```
$ crsctl query crs activeversion
```
- You must ensure that any rebalance operations on your existing Oracle ASM installation are completed before starting the upgrade or patching process.
- During the upgrade or rolling patch process, you alter the Oracle ASM instances to an upgrade state. You do not need to shut down database clients unless they are on Oracle ACFS. However, because this upgrade state limits Oracle ASM operations, you should complete the upgrade process soon after you begin. The following are the operations allowed when an Oracle ASM instance is in the upgrade state:
 - Diskgroup mounts and dismounts

- Opening, closing, resizing, or deleting database files
 - Recovering instances
 - Queries of fixed views and packages: Users are allowed to query fixed views and run anonymous PL/SQL blocks using fixed packages, such as `dbms_diskgroup`)
- You do not need to shut down database clients unless they are on Oracle ACFS.

See Also: See [Section B.10.1, "Upgrading Oracle ASM Using ASMCA"](#) for steps to upgrade Oracle ASM separately using ASMCA

B.10 Performing Rolling Upgrade of Oracle ASM

After you have completed the Oracle Clusterware portion of Oracle Grid Infrastructure 12c Release 1 (12.1) upgrade, you may need to upgrade Oracle ASM separately under the following conditions:

- If you are upgrading from a release in which Oracle ASM was in a separate Oracle home, such as Oracle ASM 10g Release 2 (10.2) or Oracle ASM 11g Release 1 (11.1)
- If the Oracle ASM portion of the Oracle Grid Infrastructure upgrade failed, or for some other reason Automatic Storage Management Configuration assistant (`asmca`) did not run.

You can use `asmca` to complete the upgrade separately, but you should do it soon after you upgrade Oracle Clusterware, as Oracle ASM management tools such as `srvctl` do not work until Oracle ASM is upgraded.

Note: ASMCA performs a rolling upgrade only if the earlier release of Oracle ASM is either 11.1.0.6 or 11.1.0.7. Otherwise, ASMCA performs a non-rolling upgrade, in which ASMCA shuts down all Oracle ASM instances on all nodes of the cluster, and then starts an Oracle ASM instance on each node from the new Oracle Grid Infrastructure home.

After you have upgraded Oracle ASM with Oracle Grid Infrastructure 12c Release 1, you can install individual patches for Oracle ASM by downloading them from the Oracle Automated Release Update site. See [Section B.9, "Restrictions and Guidelines for Upgrading and Patching Oracle ASM"](#) for more information about upgrading Oracle ASM separately using ASMCA.

B.10.1 Upgrading Oracle ASM Using ASMCA

Complete the following tasks if you must upgrade from an Oracle ASM release where Oracle ASM was installed in a separate Oracle home, or if the Oracle ASM portion of Oracle Grid Infrastructure upgrade failed to complete:

1. On the node you plan to start the upgrade, set the environment variable `ASMCA_ROLLING_UPGRADE` as true. For example:

```
$ export ASMCA_ROLLING_UPGRADE=true
```

2. From the Oracle Grid Infrastructure 12c Release 1 (12.1) home, start ASMCA. For example:

```
$ cd /u01/12.1/grid/bin
$ ./asmca
```

3. Select **Upgrade**.

ASM Configuration Assistant upgrades Oracle ASM in succession for all nodes in the cluster.

4. After you complete the upgrade, run the command to unset the `ASMCA_ROLLING_UPGRADE` environment variable.

See Also: *Oracle Database Upgrade Guide* and *Oracle Automatic Storage Management Administrator's Guide* for additional information about preparing an upgrade plan for Oracle ASM, and for starting, completing, and stopping Oracle ASM upgrades

B.11 Applying Patches to Oracle ASM

After you have upgraded Oracle ASM with Oracle Grid Infrastructure 12c Release 1, you can install individual patches for Oracle ASM by downloading them from My Oracle Support.

This section explains about Oracle ASM patches as follows:

- [About Individual \(One-Off\) Oracle ASM Patches](#)
- [About Oracle ASM Software Patch Levels](#)
- [Patching Oracle ASM to a Software Patch Level](#)

B.11.1 About Individual (One-Off) Oracle ASM Patches

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. See [Section B.11.3, "Patching Oracle ASM to a Software Patch Level"](#) for information about applying patches to Oracle ASM using the OPatch Utility.

B.11.2 About Oracle ASM Software Patch Levels

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.

- See Also:**
- [Section B.8.1, "Performing a Standard Upgrade from an Earlier Release"](#) for information about upgrading Oracle Grid Infrastructure
 - [Section B.11.3, "Patching Oracle ASM to a Software Patch Level"](#) for information about applying patches to Oracle ASM using the OPatch Utility

B.11.3 Patching Oracle ASM 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/12.1.0/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.

B.12 Updating Oracle Enterprise Manager Cloud Control Target Parameters

Because Oracle Grid Infrastructure 12c Release 1 (12.1) 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, upgrade the Enterprise Manager Cloud Control target, and then update the Enterprise Manager Agent Base Directory on each cluster member node running an agent, as described in the following sections:

- [Updating the Enterprise Manager Cloud Control Target After Upgrades](#)
- [Updating the Enterprise Manager Agent Base Directory After Upgrades](#)

B.12.1 Updating the Enterprise Manager Cloud Control Target After Upgrades

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.

B.12.2 Updating the Enterprise Manager Agent Base Directory After Upgrades

1. Navigate to the `bin` directory in the Management Agent home.

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/12.1.0.1.0`.

2. In the `/u01/app/emagent/core/12.1.0.1.0/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.

B.13 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. Complete this task using the following procedure:

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
```

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 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 12.1.0.1 installation media to remove Oracle software from an existing 11.2.0.4 Oracle home.

See Also: [Section 10.6.1, "About the Deinstallation Tool"](#)

B.14 Checking Cluster Health Monitor Repository Size After Upgrading

If you are upgrading from a prior release using IPD/OS to Oracle Grid Infrastructure then review the Cluster Health Monitor repository size (the CHM repository). Oracle recommends that you review your CHM repository needs, and enlarge the repository size if you want to maintain a larger CHM repository.

Note: Your previous IPD/OS repository is deleted when you install Oracle Grid Infrastructure, and you run the `root.sh` script on each node.

Cluster Health Monitor is not available with IBM: Linux on System z configurations.

By default, the CHM repository size is a minimum of either 1GB or 3600 seconds (1 hour). The CHM repository is one gigabyte (1 GB), regardless of the size of the cluster.

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
```

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 there is not sufficient space available, then you can move the repository to shared storage.

For example, to set the repository size to four hours:

```
$ oclumon manage -repos changeretentiontime 14400
```

B.15 Downgrading Oracle Clusterware After an Upgrade

After a successful or a failed upgrade to Oracle Clusterware 12c Release 1 (12.1), you can restore Oracle Clusterware to the previous release. This section contains the following topics:

- [About Downgrading Oracle Clusterware After an Upgrade](#)
- [Downgrading to Releases Before 11g Release 2 \(11.2.0.2\)](#)
- [Downgrading to 11g Release 1 \(11.2.0.2\) or Later Release](#)

B.15.1 About Downgrading Oracle Clusterware After an Upgrade

Downgrading Oracle Clusterware restores the Oracle Clusterware configuration to the state it was in before the Oracle Clusterware 12c Release 1 (12.1) upgrade. Any configuration changes you performed during or after the Oracle Grid Infrastructure 12c Release 1 (12.1) upgrade are removed and cannot be recovered.

In the downgrade procedures, the following variables are used:

- *first node* is the first node on which the `rootupgrade` script completed successfully.

- *non-first nodes* are all other nodes where the `rootupgrade` script completed successfully.

To restore Oracle Clusterware to the previous release, use the downgrade procedure for the release to which you want to downgrade.

Note: When downgrading after a failed upgrade, if `rootcrs.sh` does not exist on a node, then use `perl rootcrs.pl` instead of `rootcrs.sh`.

B.15.2 Downgrading to Releases Before 11g Release 2 (11.2.0.2)

To downgrade Oracle Clusterware:

1. If the `rootupgrade` script failed on a node, then downgrade the node where the upgrade failed:

```
rootcrs.sh -downgrade
```

2. On all other nodes where the `rootupgrade` script ran successfully, use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade` to stop the 12c Release 1 (12.1) resources, and shut down the Oracle Grid Infrastructure 12c Release 1 (12.1) stack.

```
rootcrs.sh -downgrade
```

3. After the `rootcrs.sh -downgrade` script has completed on all non-first nodes, on the first node use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade -lastnode`.

For example:

```
# /u01/app/12.1.0/grid/crs/install/rootcrs.sh -downgrade -lastnode
```

Note: With Oracle Grid Infrastructure 12c, you no longer need to provide the location of the previous release Grid home or release number.

Run this command from a directory that has write permissions for the Oracle Grid Infrastructure installation user.

4. On any of the cluster member nodes where the `rootcrs` script has run successfully:

- a. Log in as the Oracle Grid Infrastructure installation owner.

- b. Use the following command to start the installer, where

`/u01/app/12.1.0/grid` is the location of the new (upgraded) Grid home:

```
./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -updateNodeList
-silent
CRS=false ORACLE_HOME=/u01/app/12.1.0/grid
```

Add the flag `-cfs` if the Grid home is a shared home.

5. On any of the cluster member nodes where the `rootupgrade.sh` script has run successfully:

- a. Log in as the Oracle Grid Infrastructure installation owner (`grid`).

- b. Use the following command to start the installer, where the path you provide for the flag `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation

For example:

```
$ cd /u01/app/12.1.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
  -updateNodeList -silent
    CRS=true ORACLE_HOME=/u01/app/crs
```

- c. **For downgrades to 11.1 and earlier releases**

If you are downgrading to Oracle Clusterware 11g Release 1 (11.1) or an earlier release, then you must run `root.sh` manually from the earlier release Oracle Clusterware home to complete the downgrade after you complete step b.

OUI prompts you to run `root.sh` manually from the earlier release Oracle Clusterware installation home in sequence on each member node of the cluster to complete the downgrade. After you complete this task, downgrade is completed.

Running `root.sh` from the earlier release Oracle Clusterware installation home restarts the Oracle Clusterware stack, starts up all the resources previously registered with Oracle Clusterware in the earlier release, and configures the old initialization scripts to run the earlier release Oracle Clusterware stack.

After completing the downgrade, update the entry for Oracle ASM instance in the `oratab` file (`/etc/oratab` or `/var/opt/oracle/oratab`) on every node in the cluster as follows:

```
+ASM<instance#>:<RAC-ASM home>:N
```

B.15.3 Downgrading to 11g Release 1 (11.2.0.2) or Later Release

Follow these steps to downgrade Oracle Grid Infrastructure:

1. On all remote nodes, use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade` to stop the 12c Release 1 (12.1) resources, and shut down the Oracle Grid Infrastructure 12c Release 1 (12.1) stack.

```
# /u01/app/12.1.0/grid/crs/install/rootcrs.sh -downgrade
```

2. After the `rootcrs.sh -downgrade` script has completed on all remote nodes, on the local node use the command syntax `Grid_home/crs/install/rootcrs.sh -downgrade -lastnode`

For example:

```
# /u01/app/12.1.0/grid/crs/install/rootcrs.sh -downgrade -lastnode
```

Note: Starting with Oracle Grid Infrastructure 12c Release 1 (12.1), you no longer need to provide the location of the earlier release Grid home or earlier release number.

Run this command from a directory that has write permissions for the Oracle Grid Infrastructure installation user.

3. On any of the cluster member nodes where the `rootupgrade.sh` script has run successfully:

- a. Log in as the Oracle Grid Infrastructure installation owner.
- b. Use the following command to start the installer, where `/u01/app/12.1.0/grid` is the location of the new (upgraded) Grid home:

```
$ cd /u01/app/12.1.0/grid/oui/bin
./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs -updateNodeList
-silent CRS=false ORACLE_HOME=/u01/app/12.1.0/grid
```

Add the flag `-cfs` if the Grid home is a shared home.

4. On any of the cluster member nodes where the `rootupgrade` script has run successfully:

- a. 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 flag `ORACLE_HOME` is the location of the home directory from the earlier Oracle Clusterware installation

For example:

```
$ cd /u01/app/12.1.0/grid/oui/bin
$ ./runInstaller -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=true ORACLE_HOME=/u01/app/crs
```

- c. **For downgrades to 11.2.0.2**

If you are downgrading to Oracle Clusterware 11g Release 1 (11.2.0.2), then you must start the Oracle Clusterware stack manually after you complete step b.

On each node, start Oracle Clusterware 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
```

5. **For downgrades to 12.1.0.1**

If you are downgrading to Oracle Grid Infrastructure 12c Release 1 (12.1.0.1), then run the following commands to configure the Grid Management Database:

- a. Start the 12.1.0.1 Oracle Clusterware stack on all nodes.
- b. On any node, remove the `MGMTDB` resource as follows:

```
12101_Grid_home/bin/srvctl remove mgmtdb
```

- c. 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 _mgmtdb -storageType ASM
-diskGroupName ASM_DG_NAME
-datafileJarLocation 12101_grid_home/assistants/dbca/templates
-characterSet AL32UTF8 -autoGeneratePasswords
```

- d. Configure the Management Database by running the Configuration Assistant from the location `12101_Grid_home/bin/mgmtca`.

Installing and Configuring Oracle Database Using Response Files

This appendix describes how to install and configure Oracle products using response files. It includes information about the following topics:

- [How Response Files Work](#)
- [Preparing a Response File](#)
- [Running the Installer Using a Response File](#)
- [Running Net Configuration Assistant Using a Response File](#)
- [Postinstallation Configuration Using a Response File](#)

C.1 How Response Files Work

When you start the installer, you can use a response file to automate the installation and configuration of Oracle software, either fully or partially. The installer 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, supply the appropriate value for the `ORACLE_HOME` variable:

```
ORACLE_HOME="OraDBHome1"
```

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 "ORACLE_HOME=OraDBHome1" ...
```

Note that if you use a response file, you are required to edit the response file manually to enter values for passwords. To protect system security, you cannot save passwords in a response file.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for more information about response files

C.1.1 Reasons for Using Silent Mode or Response File Mode

The following table provides use cases for running the installer in silent mode or response file mode.

Mode	Uses
Silent	<p>Use silent mode to do 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 multiple systems, providing default answers to some, but not all of the installer prompts.</p> <p>In response file mode, all the installer screens are displayed, but defaults for the fields in these screens are provided by the response file. You have to provide information for the fields in screens where you have not provided values in the response file.</p>

C.1.2 General Procedure for Using Response Files

The following are the 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. Prepare a response file.
2. Run the installer in silent or response file mode.
3. If you completed a software-only installation, then run Net Configuration Assistant and Database Configuration Assistant in silent or response file mode.

These steps are described in the following sections.

C.2 Preparing a Response File

This section describes the following methods to prepare a response file for use during silent mode or response file mode installations:

- [Editing a Response File Template](#)
- [Recording a Response File](#)

C.2.1 Editing a Response File Template

Oracle provides response file templates for each product and installation type, and for each configuration tool. These files are located at `database/response` directory on the installation media.

Note: If you copied the software to a hard disk, then the response files are located in the directory `/response`.

Table C-1 lists the response files provided with this software:

Table C-1 Response Files for Oracle Database

Response File	Description
<code>db_install.rsp</code>	Silent installation of Oracle Database 11g
<code>dbca.rsp</code>	Silent installation of Database Configuration Assistant
<code>netca.rsp</code>	Silent installation of Oracle Net Configuration Assistant

Table C-2 Response files for Oracle Grid Infrastructure

Response File	Description
<code>grid_install.rsp</code>	Silent installation 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 /directory_path/response/response_file.rsp local_directory
```

In this example, `directory_path` is the path to the database directory on the installation media. If you have copied the software to a hard drive, then you can edit the file in the `response` directory.

2. Open the response file in a text editor:

```
$ vi /local_dir/response_file.rsp
```

See Also: *Oracle Universal Installer NextGen Installation Guide* for detailed information on creating response files

3. Review parameters in the response file, and provide values for your cluster.

Note: The installer or configuration assistant fails if you do not correctly configure the response file.

4. Change the permissions on the file to 600:

```
$ chmod 600 /local_dir/response_file.rsp
```

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). Ensure that only the Oracle software owner user can view or modify response files or consider deleting them after the installation succeeds.

C.2.2 Recording a Response File

You can use the installer in interactive mode to record a response file, which you can edit and then use to complete silent mode or response file mode installations. This method is useful for custom or software-only installations.

Starting with Oracle Database 11g Release 2 (11.2), 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 it 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: Oracle Universal Installer (OUI) does not record passwords in the response file.

To record a response file:

1. Complete preinstallation tasks as for a normal installation.
2. Ensure that the Oracle Grid Infrastructure software owner user (typically `grid`) has permissions to create or write to the Grid 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** and specify a file name and location to save the values for the response file, and click **Save**.
 - b. Click **Finish** to create the response file and continue with the installation.

Click **Save Response File** and **Cancel** if you only want to create the response file but not continue with the installation. The installation will stop, but the settings you have entered will be recorded in the response file.

5. Before you use the saved response file on another system, edit the file and make any required changes.

Review parameters in the response file, and provide values for your cluster.

C.3 Running the Installer Using a Response File

Run Oracle Universal Installer at the command line, specifying the response file you created. The Oracle Universal Installer executable, `runInstaller`, provides several options. For help information on the full set of these options, run the `runInstaller` command with the `-help` option. For example:

```
$ directory_path/runInstaller -help
```

The help information appears in a window.

To run the installer using a response file:

1. Complete the preinstallation tasks as 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:

```
$ /directory_path/runInstaller [-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 `oraInstRoot.sh` and `root.sh` scripts. For example:

```
$ su root  
password:
```

```
# /oracle_home_path/orainstRoot.sh
```

C.4 Running Net Configuration Assistant Using a Response File

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 on the DVD.

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.

2. Open the response file in a text editor:

```
$ vi /local_dir/netca.rsp
```

3. Review parameters in the response file, and provide values for your cluster.

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 operating system `ORACLE_HOME` environment variable for that owner 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 runs Net Configuration Assistant in silent mode.
- `local_dir` is the full path of the directory where you copied the `netca.rsp` response file template.

C.5 Postinstallation Configuration Using a Response File

Use the following sections to create and run a response file configuration after installing Oracle software.

C.5.1 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 creating and 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 it in a secure location. In addition, if you have to stop an installation to fix an error, you can run the configuration assistants using `configToolAllCommands` and a password response file.

The `configToolAllCommands` password response file consists of the following syntax options:

- *internal_component_name* is the name of the component that the configuration assistant configures
- *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.assistants.asm | S_ASMPASSWORD=welcome
```

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.

C.5.2 Running Postinstallation Configuration Using a Response File

To run configuration assistants with the `configToolAllCommands` script:

1. Create a response file using the syntax *filename.properties*. For example:

```
$ touch cfgrsp.properties
```

2. Open the file with a text editor, and cut and paste the password template, modifying as needed.

Example C-1 Password response file for Oracle Grid Infrastructure installation for a cluster

Oracle Grid Infrastructure requires passwords for Oracle Automatic Storage Management Configuration Assistant (ASMCA), and for Intelligent Platform Management Interface Configuration Assistant (IPMICA) if you have a BMC card and you want to enable this feature. Provide the following response file:

```
oracle.assistants.asm | S_ASMPASSWORD=password
oracle.assistants.asm | S_ASMMONITORPASSWORD=password
oracle.crs | S_BMCPASSWORD=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.

Note: If you are upgrading Oracle ASM 11g Release 1 or earlier releases, then you only need to provide the input field for `oracle.assistants.asm|S_ASMMONITORPASSWORD`.

Example C-2 Password response file for Oracle Real Application Clusters

Oracle Database configuration requires the SYS, SYSTEM, and DBSNMP passwords for use with Database Configuration Assistant (DBCA). Providing a string for the `S_ASMSNMPPASSWORD` variable is necessary only if the database is using Oracle ASM for storage. Also, providing a string for the `S_PDBADMINPASSWORD` variable is necessary only if you create a multitenant container database (CDB) with one or more pluggable databases (PDBs). Also, if you selected to configure Oracle Enterprise Manager Cloud Control, then you must provide the password for the Oracle software installation owner for the `S_EMADMINPASSWORD` variable, similar to the following example, where the phrase *password* represents the password string:

```
oracle.assistants.server|S_SYSPASSWORD=password
oracle.assistants.server|S_SYSTEMPASSWORD=password
oracle.assistants.server|S_DBSNMPPASSWORD=password
oracle.assistants.server|S_PDBADMINPASSWORD=password
oracle.assistants.server|S_EMADMINPASSWORD=password
oracle.assistants.server|S_ASMSNMPPASSWORD=password
```

If you do not want to enable Oracle Enterprise Manager for Oracle ASM, then leave those password fields blank.

3. Change permissions to secure the file. For example:

```
$ ls -al cfgrsp.properties
-rw----- 1 oracle oinstall 0 Apr 30 17:30 cfgrsp
```

4. Change directory to `$ORACLE_HOME/cfgtoollogs`, and run the configuration script using the following syntax:

```
configToolAllCommands RESPONSE_FILE=/path/name.properties
```

for example:

```
$ ./configToolAllCommands RESPONSE_FILE=/home/oracle/cfgrsp.properties
```

Configuring Large Memory Optimization

This appendix provides information for configuring memory optimization with large page tables on the Linux operating system, using Hugepages. It contains the following topics:

- [Overview of HugePages](#)
- [Restrictions for HugePage Configurations](#)
- [Disabling Transparent HugePages](#)

D.1 Overview of HugePages

You can choose to configure HugePages. For some uses, HugePages can provide enhanced performance. However, this feature is an advanced configuration option. It is not a requirement for Oracle Real Application Clusters (Oracle RAC).

The following is an overview of HugePages. It does not provide RPM or configuration information. The tasks you must perform for configuration depend on kernel distribution and hardware on your system. If you decide to configure your cluster nodes to use HugePages, then refer to your distribution documentation and to Oracle Technology Network and My Oracle Support for further information.

D.1.1 What HugePages Provides

HugePages is a feature integrated into the Linux kernel with release 2.6. It is a method to have larger pages where it is useful for working with very large memory. It can be useful for both 32-bit and 64-bit configurations. HugePage sizes vary from 2 MB to 256 MB, depending on the kernel version and the hardware architecture. For Oracle Database, using HugePages reduces the operating system maintenance of page states, and increases TLB (Translation Lookaside Buffer) hit ratios.

Without HugePages, the operating system keeps each 4 KB of memory as a page. When that memory is allocated to the SGA, the lifecycle of that page (dirty, free, mapped to a process, and so on) must be kept up to date by the operating system kernel.

With HugePages, the operating system page table (virtual memory to physical memory mapping) is smaller, because each page table entry is pointing to pages from 2 MB to 256 MB. Also, the kernel has fewer pages whose lifecycle must be monitored.

For example, if you use HugePages with 64-bit hardware, and you want to map 256 MB of memory, you may need one page table entry (PTE). If you do not use HugePages, and you want to map 256 MB of memory, then you must have $256 \text{ MB} * 1024 \text{ KB} / 4 \text{ KB} = 65536 \text{ PTEs}$.

D.2 Restrictions for HugePage Configurations

The HugePages feature allocates non-swappable memory for large page tables using memory-mapped files. If you enable HugePages, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.

To use HugePages, you must configure Grub to allocate memory for HugePages during system startup. After paging space is reserved, HugePages can be used as needed. However, if the space they require is not reserved in memory during system startup, then a HugePages allocation may fail.

You must also ensure that both `MEMORY_TARGET` and `MEMORY_MAX_TARGET` initialization parameters are unset (for example, use the command `ALTER SYSTEM RESET`) for the database instance.

HugePages memory is not subject to allocation or release after system startup, unless a system administrator changes the HugePages configuration by modifying the number of pages available, or the pool size.

See Also: *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for more information about HugePages

D.3 Disabling Transparent HugePages

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

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.

See Also: *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for information about HugePages

To check if Transparent HugePages memory 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 memory being used as the `[always]` flag is enabled.

```
[always] never
```

Note: If Transparent HugePages is removed from the kernel, then the `/sys/kernel/mm/transparent_hugepage` or `/sys/kernel/mm/redhat_transparent_hugepage` files do not exist.

To disable Transparent HugePages, perform the following steps:

1. 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
```

2. Restart the system to make the changes permanent.

Oracle Grid Infrastructure for a Cluster Installation Concepts

This appendix provides an overview of concepts and terms that may be necessary to carry out installation.

This appendix contains the following sections:

- [Understanding Preinstallation Configuration](#)
- [Understanding Network Addresses](#)
- [Understanding Network Time Requirements](#)
- [Understanding Oracle Flex Clusters and Oracle ASM Flex Clusters](#)
- [Understanding Storage Configuration](#)
- [Understanding Out-of-Place Upgrade](#)

E.1 Understanding Preinstallation Configuration

This section reviews concepts about Oracle Grid Infrastructure for a cluster preinstallation tasks. It contains the following sections:

- [Optimal Flexible Architecture Guidelines for Oracle Grid Infrastructure](#)
- [Oracle Grid Infrastructure for a Cluster and Oracle Restart Differences](#)
- [Understanding the Oracle Inventory Group](#)
- [Understanding the Oracle Inventory Directory](#)
- [Understanding the Oracle Base directory](#)
- [Understanding the Oracle Home for Oracle Grid Infrastructure Software](#)
- [Location of Oracle Base and Oracle Grid Infrastructure Software Directories](#)

E.1.1 Optimal Flexible Architecture Guidelines for Oracle Grid Infrastructure

For installations with Oracle Grid Infrastructure only, Oracle recommends that you create an Oracle base and Grid home path compliant with Oracle Optimal Flexible Architecture (OFA) guidelines, so that Oracle Universal Installer (OUI) can select that directory during installation. For OUI to recognize the path as an Oracle software path, it must be in the form `u[0-9][1-9]/app`.

The OFA path for an Oracle base is `u[0-9][1-9]/app/user`, where *user* is the name of the Oracle software installation owner account.

The OFA path for an Oracle Grid Infrastructure Oracle home is `u[0-9][1-9]/app/release/grid` where *release* is the three-digit Oracle Grid Infrastructure release (for example, 12.1.0).

When OUI finds an OFA-compliant software path (`u[0-9][1-9]/app`), it creates the Oracle Grid Infrastructure Grid home and Oracle Inventory (`oraInventory`) directories for you. For example, the path `/u01/app` and `/u89/app` are OFA-compliant paths.

The Oracle Grid Infrastructure home must be in a path that is different from the Grid home for the Oracle Grid Infrastructure installation owner. If you create an Oracle Grid Infrastructure base path manually, then ensure that it is in a separate path specific for this release, and not under an existing Oracle base path.

Note: If you choose to create an Oracle Grid Infrastructure home manually, then do not create the Oracle Grid Infrastructure home for a cluster under either the Oracle Grid Infrastructure installation owner (`grid`) Oracle base or the Oracle Database installation owner (`oracle`) Oracle base. Creating an Oracle Clusterware installation in an Oracle base directory will cause succeeding Oracle installations to fail.

Oracle Grid Infrastructure homes can be placed in a local home on servers, even if your existing Oracle Clusterware home from a prior release is in a shared location.

E.1.2 Oracle Grid Infrastructure for a Cluster and Oracle Restart Differences

Requirements for Oracle Grid Infrastructure for a cluster are different from Oracle Grid Infrastructure on a single instance in an Oracle Restart configuration.

See Also: *Oracle Database Installation Guide* for information about Oracle Restart requirements

E.1.3 Understanding the Oracle Inventory Group

You must have a group whose members are given access to write to the Oracle Inventory (`oraInventory`) directory, which is the central inventory record of all Oracle software installations on a server. Members of this group have write privileges 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. By default, this group is called `oinstall`. The Oracle Inventory group must be the primary group for Oracle software installation owners.

The `oraInventory` directory contains the following:

- A registry of the Oracle home directories (Oracle Grid Infrastructure and Oracle Database) on the system.
- Installation logs and trace files from installations of Oracle software. These files are also copied to the respective Oracle homes for future reference.
- Other metadata inventory information regarding Oracle installations are stored in the individual Oracle home inventory directories, and are separate from the central inventory.

You can configure one group to be the access control group for the Oracle Inventory, for database administrators (OSDBA), and for all other access control groups used by Oracle software for operating system authentication. However, if you use one group to provide operating system authentication for all system privileges, then this group

must be the primary group for all users to whom you want to grant administrative system privileges.

Note: If Oracle software is already installed on the system, then the existing Oracle Inventory group must be the primary group of the operating system user (`oracle` or `grid`) that you use to install Oracle Grid Infrastructure. See [Section 6.1.1, "Determining If the Oracle Inventory and Oracle Inventory Group Exists"](#) to identify an existing Oracle Inventory group.

E.1.4 Understanding the Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) is the central inventory location for all Oracle software installed on a server. Each cluster member node has its own central inventory file. You cannot have a shared Oracle Inventory directory, because it is used to point to the installed Oracle homes for all Oracle software installed on a node.

The first time you install Oracle software on a system, you are prompted to provide an `oraInventory` directory path.

By default, if an `oraInventory` group does not exist, then the installer lists the primary group of the installation owner for the Oracle Grid Infrastructure for a cluster software as the `oraInventory` group. Ensure that this group is available as a primary group for all planned Oracle software installation owners.

The primary group of all Oracle installation owners should be the Oracle Inventory Group (`oinstall`), whose members are granted the `OINSTALL` system privileges to write to the central Oracle Inventory for a server, to write log files, and other privileges.

Note: Group and user IDs must be identical on all nodes in the cluster. Check to make sure 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.

If the primary group of an installation owner is the user home directory (for example, `/home/oracle`), then the Oracle Inventory is placed in the installation owner's home directory. This placement can cause permission errors during subsequent installations with multiple Oracle software owners. For that reason, Oracle recommends that you do not accept this option, and instead use an OFA-compliant path.

If you set an Oracle base variable to a path such as `/u01/app/grid` or `/u01/app/oracle`, then the Oracle Inventory is defaulted to the path `u01/app/oraInventory` using correct permissions to allow all Oracle installation owners to write to this central inventory directory.

By default, the Oracle Inventory directory is not installed under the Oracle base directory for the installation owner. This is because all Oracle software installations share a common Oracle Inventory, so there is only one Oracle Inventory for all users, whereas there is a separate Oracle base for each user.

E.1.5 Understanding the Oracle Base directory

During installation, you are prompted to specify an Oracle base location, which is owned by the user performing the installation. The Oracle base directory is where log

files specific to the user are placed. You can choose a location with an existing Oracle home, or choose another directory location that does not have the structure for an Oracle base directory.

Using the Oracle base directory path helps to facilitate the organization of Oracle installations, and helps to ensure that installations of multiple databases maintain an Optimal Flexible Architecture (OFA) configuration.

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 as an Oracle software path, it must be in the form `u[0-9][1-9]/app`, and it must be writable by any member of the `oraInventory` (`oinstall`) group. The OFA path for the Oracle base is `u[0-9][1-9]/app/user`, where *user* is the name of the software installation owner. For example:

```
/u01/app/grid
```

Because you can have only one Oracle Grid Infrastructure installation on a cluster, and all upgrades are out-of-place upgrades, Oracle recommends that you create an Oracle base for the grid infrastructure software owner (`grid`), and create an Oracle home for the Oracle Grid Infrastructure binaries using the release number of that installation.

E.1.6 Understanding the Oracle Home for Oracle Grid Infrastructure Software

The Oracle home for Oracle Grid Infrastructure software (Grid home) should be in a path in the format `u[0-9][1-9]/app/release/grid`, where *release* is the release number of the Oracle Grid Infrastructure software. For example:

```
/u01/app/12.1.0/grid
```

During installation, ownership of the path to the Grid home is changed to `root`. 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.

Ensure that the directory path you provide for the Oracle Grid Infrastructure software location (Grid home) complies with the following requirements:

- If you create the path before installation, then it should be owned by the installation owner of Oracle Grid Infrastructure (typically `oracle` for a single installation owner for all Oracle software, or `grid` for role-based Oracle installation owners), and set to 775 permissions.
- It should be created in a path outside existing Oracle homes, including Oracle Clusterware homes.
- It should not be located in a user home directory.
- It must not be the same location as the Oracle base for the Oracle Grid Infrastructure installation owner (`grid`), or the Oracle base of any other Oracle installation owner (for example, `/u01/app/oracle`).
- It should be created either as a subdirectory in a path where all files can be owned by `root`, or in a unique path.

Oracle recommends that you install Oracle Grid Infrastructure binaries on local homes, rather than using a shared home on shared storage.

E.1.7 Location of Oracle Base and Oracle Grid Infrastructure Software Directories

Even if you do not use the same software owner to install Grid Infrastructure (Oracle Clusterware and Oracle ASM) and Oracle Database, be aware that running the `root.sh` script during the Oracle Grid Infrastructure installation changes ownership of the home directory where clusterware binaries are placed to `root`, and all ancestor directories to the root level (`/`) are also changed to `root`. For this reason, the Oracle Grid Infrastructure for a cluster home cannot be in the same location as other Oracle software.

However, Oracle Restart can be in the same location as other Oracle software.

See Also: *Oracle Database Installation Guide* for your platform for more information about Oracle Restart

E.2 Understanding Network Addresses

During installation, you are asked to identify the planned use for each network interface that 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 Flex ASM cluster to use. Public and virtual IP addresses are configured on public interfaces. Private addresses are configured on private interfaces.

See the following sections for detailed information about each address type:

- [About the Public IP Address](#)
- [About the Private IP Address](#)
- [About the Virtual IP Address](#)
- [About the Grid Naming Service \(GNS\) Virtual IP Address](#)
- [About the SCAN for Oracle Grid Infrastructure Installations](#)

E.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. 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 will appear to be down until it is removed from the cluster.

E.2.2 About the Private IP Address

Oracle Clusterware uses interfaces marked as private for internode communication. Each cluster node needs to have an interface that you identify during installation as a private interface. Private interfaces need to have addresses configured for the interface itself, but no additional configuration is required. Oracle Clusterware uses 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.

After installation, if you modify interconnects on Oracle RAC with the `CLUSTER_INTERCONNECTS` initialization parameter, then you must change it to a private IP address, on a subnet that is not used with a public IP address. 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, as this can block interconnect traffic.

E.2.3 About the Virtual IP Address

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 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 *hostname-vip*. For example: `myclstr2-vip`.

The virtual IP (VIP) address is registered in the 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

E.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 will serve. The set of IP addresses is provided to the cluster through DHCP, which must be available on the public network for the cluster.

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about Grid Naming Service

E.2.5 About the SCAN for Oracle Grid Infrastructure Installations

Oracle Database clients connect to Oracle Real Application Clusters database using SCANs. 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 works by being able to resolve to multiple IP addresses in the cluster handling 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. They listen on the SCAN IP addresses provided on nodes for the SCAN IP addresses. 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 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.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.grid.example.com`.

Clients configured to use IP addresses for Oracle Database releases before Oracle Database 11g Release 2 can continue to use their existing connection addresses; using SCANS 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 version 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 with the SCAN as `HOST=SCAN`.

The SCAN is optional for most deployments. However, clients using Oracle Database 11g Release 2 and later policy-managed databases using server pools should access the database using the SCAN. This is because policy-managed databases can run on different servers at different times, so connecting to a particular node virtual IP address for a policy-managed database is not possible.

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

Note: 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.
-
-

Identify public and private interfaces. OUI 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.

E.3 Understanding Network Time Requirements

Oracle Clusterware 12c Release 1 (12.1) is automatically configured with Cluster Time Synchronization Service (CTSS). This service provides automatic synchronization of all cluster nodes using the optimal synchronization strategy for the type of cluster you deploy. If you have existing cluster time synchronization service, such as NTP, then it starts in an observer mode. Otherwise, it starts in an active mode to ensure that time is synchronized between cluster nodes. CTSS does not cause compatibility issues.

The CTSS module is installed as a part of Oracle Grid Infrastructure installation. CTSS daemons are started up by the OHAS daemon (`ohasd`), and do not require a command-line interface.

E.4 Understanding Oracle Flex Clusters and Oracle ASM Flex Clusters

Oracle Grid Infrastructure installed in an Oracle Flex Cluster configuration is a scalable, dynamic, robust network of nodes. 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.

Oracle Flex Cluster Hub Nodes are similar to Oracle Grid Infrastructure nodes in a standard configuration: they are tightly connected, and have direct access to shared storage.

Leaf Nodes are different from standard Oracle Grid Infrastructure nodes, in that they do not require direct access to shared storage. 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.

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Flex Cluster deployments
- *Oracle Automatic Storage Management Administrator's Guide* for information about Oracle Flex ASM

E.5 Understanding Storage Configuration

[Understanding Oracle Automatic Storage Management Cluster File System](#)

[About Migrating Existing Oracle ASM Instances](#)

[Standalone Oracle ASM Installations to Clustered Installation Conversions](#)

E.5.1 Understanding Oracle Automatic Storage Management Cluster File System

Oracle Automatic Storage Management has been extended to include a general purpose file system, called Oracle Automatic Storage Management Cluster File System (Oracle ACFS). Oracle ACFS is a new multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of the Oracle Database. Files supported by Oracle ACFS include application binaries and application reports. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Automatic Storage Management Cluster File System (ACFS) can provide optimized storage for all Oracle files, including Oracle Database binaries. It can also store other application files. However, it cannot be used for Oracle Clusterware binaries.

See Also : *Oracle Automatic Storage Management Administrator's Guide* for more information about ACFS

E.5.2 About Migrating Existing Oracle ASM Instances

If you have an Oracle ASM installation from a prior release installed on your server, or in an existing Oracle Clusterware installation, then you can use Oracle Automatic Storage Management Configuration Assistant (ASMCA, located in the path *Grid_home/bin*) to upgrade the existing Oracle ASM instance to Oracle ASM 12c Release 1 (12.1), and subsequently configure failure groups, Oracle ASM volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

Note: You must first shut down all database instances and applications on the node with the existing Oracle ASM instance before upgrading it.

During installation, if you chose to use Oracle ASM and ASMCA detects that there is a prior Oracle ASM version installed in another home, then after installing the Oracle ASM 12c Release 1 (12.1) binaries, you can start ASMCA to upgrade the existing Oracle ASM instance. You can then choose to configure an Oracle ACFS deployment

by creating Oracle ASM volumes and using the upgraded Oracle ASM to create the Oracle ACFS.

On an existing Oracle Clusterware or Oracle RAC installation, if the earlier version of Oracle ASM instances on all nodes is Oracle ASM 11g Release 1 (11.1), then you are provided with the option to perform a rolling upgrade of Oracle ASM instances. If the earlier version of Oracle ASM instances on an Oracle RAC installation are from an Oracle ASM release before Oracle ASM 11g Release 1 (11.1), then rolling upgrades cannot be performed. Oracle ASM is then upgraded on all nodes to 12c Release 1 (12.1).

E.5.3 Standalone Oracle ASM Installations to Clustered Installation Conversions

If you have existing standalone Oracle ASM installations on one or more nodes that are member nodes of the cluster, then OUI proceeds to install Oracle Grid Infrastructure for a cluster.

If you place Oracle Clusterware files (OCR and voting files) on Oracle ASM, then ASMCA is started at the end of the clusterware installation, and provides prompts for you to migrate and upgrade the Oracle ASM instance on the local node, so that you have an Oracle ASM 12c Release 1 (12.1) installation.

On remote nodes, ASMCA identifies any standalone Oracle ASM instances that are running, and prompts you to shut down those Oracle ASM instances, and any database instances that use them. ASMCA then extends clustered Oracle ASM instances to all nodes in the cluster. However, disk group names on the cluster-enabled Oracle ASM instances must be different from existing standalone disk group names.

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

E.6 Understanding Out-of-Place Upgrade

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.

See Also: [Appendix B, "How to Upgrade to Oracle Grid Infrastructure 12c Release 1"](#) for instructions on completing rolling upgrades

How to Complete Preinstallation Tasks Manually

This appendix provides instructions to complete configuration tasks manually that Cluster Verification Utility (CVU) and Oracle Universal Installer (OUI) normally complete during installation using Fixup scripts. Use this appendix as a guide if you cannot use Fixup scripts.

This appendix contains the following information:

- [Configuring SSH Manually on All Cluster Nodes](#)
- [Configuring Kernel Parameters](#)
- [Setting UDP and TCP Kernel Parameters Manually](#)
- [Configuring Storage Paths and Disk Devices](#)
- [Checking OCFS2 Version Manually](#)

F.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.

In the examples that follow, the Oracle software owner listed is the `grid` user.

If SSH is not available, then OUI attempts to use `rsh` and `rcp` instead. However, these services are disabled by default on most Linux systems.

Note: The supported version of SSH for Linux distributions is OpenSSH.

This section contains the following:

- [Checking Existing SSH Configuration on the System](#)
- [Configuring SSH on Cluster Nodes](#)

- [Enabling SSH User Equivalency on Cluster Nodes](#)

F.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.

F.1.2 Configuring SSH on Cluster Nodes

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.

You must configure SSH separately for each Oracle software installation owner that you intend to use for installation.

To configure SSH, complete the following:

F.1.2.1 Create SSH Directory, and Create SSH Keys On Each 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` and `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:

```
$ id
uid=1100(grid) gid=1000(oinstall) groups=1000(oinstall)
1100(grid,asmadmin,asmdba)
$ id grid
uid=1100(grid) gid=1000(oinstall) groups=1000(oinstall),
1100(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: SSH configuration will fail 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**).

Note: SSH with passphrase is not supported for Oracle Clusterware 11g Release 2 and later releases.

This command writes the DSA public key to the `~/.ssh/id_dsa.pub` file and the private key to the `~/.ssh/id_dsa` file.

Never distribute the private key to anyone not authorized to perform Oracle software installations.

5. Repeat steps 1 through 4 on each node that you intend to make a member of the cluster, using the DSA key.

F.1.2.2 Add All Keys to a Common `authorized_keys` File

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/
```

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.

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
```

Note: 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.

F.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 the order listed. In this example, the Oracle Grid Infrastructure software owner is named `grid`:

- 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
.
.
.
```

For example:

```
[grid@node1 grid]$ ssh node1 date
The authenticity of host 'node1 (xxx.xxx.100.101)' can't be established.
DSA key fingerprint is 7z:60:60:zz:48:48:z1:a0:f7:4e.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'node1,xxx.xxx.100.101' (DSA) to the list of
known hosts.
Mon Dec 4 11:08:13 PST 2006
[grid@node1 grid]$ ssh node1.example.com date
The authenticity of host 'node1.example.com (xxx.xxx.100.101)' can't be
established.
DSA key fingerprint is 7z:60:60:zz:48:48:z1:a0:f7:4e.
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added 'node1.example.com,xxx.xxx.100.101' (DSA) to the
list of known hosts.
Mon Dec 4 11:08:13 PST 2006
[grid@node1 grid]$ ssh node2 date
Mon Dec 4 11:08:35 PST 2006
.
.
.
```

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, proceed to [Section 6.2.4, "Setting Remote Display and X11 Forwarding Configuration."](#)

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.

F.2 Configuring Kernel Parameters

This section contains the following:

- [Minimum Parameter Settings for Installation](#)
- [Additional Parameter and Kernel Settings for SUSE Linux Enterprise Server](#)

Note: The kernel parameter and shell limit values shown in the following section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. Refer to your operating system documentation for more information about tuning kernel parameters.

F.2.1 Minimum Parameter Settings for Installation

During installation, or when you run the Cluster Verification Utility (`cluvfy`) with the flag `-fixup`, a fixup script is generated. This script updates required kernel parameters if necessary to minimum values.

If you cannot use the fixup scripts, then review [Table F-1](#) to set values manually:

Table F-1 Minimum Operating System Parameter Settings for Installation on Linux

Parameter	Value	File
semmsl	250	/proc/sys/kernel/sem
semms	32000	
semopm	100	
semmni	128	
shmall	40 percent of the size of physical memory in pages. 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
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
panic_on_oops	1	/proc/sys/kernel/panic_on_oops

Note: If you intend to install Oracle Databases or Oracle RAC databases on the cluster, be aware that the size of the `/dev/shm` mount area on each server must be greater than the system global area (SGA) and the program global area (PGA) of the databases on the servers. Review expected SGA and PGA sizes with database administrators to ensure that you do not have to increase `/dev/shm` after databases are installed on the cluster.

F.2.2 Additional Parameter and Kernel Settings for SUSE Linux Enterprise Server

On SUSE Linux Enterprise Server systems only, complete the following steps as needed:

1. Enter the following command to cause 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`. Doing this grants members of `oinstall` a group permission to create shared memory segments.

For example, where the `oinstall` group GID is 1000:

```
# echo 1000 > /proc/sys/vm/hugetlb_shm_group
```

After running this command, 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=1000
```

Note: Only one group can be defined as the `vm.hugetlb_shm_group`.

3. Repeat steps 1 through 3 on all other nodes in the cluster.

F.3 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 manually 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 OUI 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 detailed information about how to automate this ephemeral port range alteration on system restarts.

F.4 Configuring Storage Paths and Disk Devices

For persistent device naming, you can configure ASMLIB or set `udev` rules.

This section consists of the following:

- [Configuring Storage Device Path Persistence Using Oracle ASMLIB](#)

- [Configuring Disk Devices Manually for Oracle ASM](#)

F.4.1 Configuring Storage Device Path Persistence Using Oracle ASMLIB

Oracle recommends that you use Oracle ASM Filter Driver (ASMFD) to maintain device persistence. However, you can choose to use ASMLIB for device persistence.

Review the following section to configure Oracle ASMLIB:

- [About Oracle ASM with Oracle ASMLIB](#)
- [Configuring Oracle ASMLIB to Maintain Block Devices](#)
- [Configuring Oracle ASMLIB for Multipath Disks](#)
- [Deinstalling Oracle ASMLIB](#)

Note: Oracle ASMLIB is not supported on IBM:Linux on System z.

F.4.1.1 About Oracle ASM with Oracle ASMLIB

The Oracle Automatic Storage Management (Oracle ASM) library driver (ASMLIB) simplifies the configuration and management of block disk devices by eliminating the need to rebind block disk devices used with Oracle ASM each time the system is restarted.

With ASMLIB, you define the range of disks you want to have made available as Oracle ASM disks. ASMLIB maintains permissions and disk labels that are persistent on the storage device, so that label is available even after an operating system upgrade. You can update storage paths on all cluster member nodes by running one `oracleasm` command on each node, without the need to modify the `udev` file manually to provide permissions and path persistence.

Note: If you configure disks using ASMLIB, then you must change the disk discovery string to `ORCL:*`. If the disk string is set to `ORCL:*`, or is left empty (`""`), then the installer discovers these disks.

F.4.1.2 Configuring Oracle ASMLIB to Maintain Block Devices

To use the Oracle Automatic Storage Management Library Driver (ASMLIB) to configure Oracle ASM devices, complete the following tasks.

- [Installing and Configuring the Oracle ASM Library Driver Software](#)
- [Configuring Disk Devices to Use Oracle ASM Library Driver on x86 Systems](#)
- [Administering the Oracle ASM Library Driver and Disks](#)

Note: To create a database during the installation using the Oracle ASM library driver, you must choose an installation method that runs ASMCA in interactive mode. You must also change the default disk discovery string to `ORCL:*`.

F.4.1.2.1 Installing and Configuring the Oracle ASM Library Driver Software ASMLIB is already included with Oracle Linux packages, and with SUSE Linux Enterprise Server. If you are a member of the Unbreakable Linux Network, then you can install the 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 ASMLIB driver software manually, follow these steps:

1. Enter the following command to determine the kernel version and architecture of the system:

```
# uname -rm
```

2. Download the required ASMLIB 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 ASMLIB on Red Hat Enterprise Linux 5 Advanced Server. ASMLIB is already included with SUSE Linux Enterprise Server distributions.

See Also: My Oracle Support Note 1089399.1 for information about ASMLIB support with Red Hat distributions:

<https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1089399.1>

3. Switch user to the root user:

```
$ su -
```

4. Install the following packages in sequence, where *version* is the version of the ASMLIB 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
oracleasmlib-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 \
    oracleasmlib-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 \
    oracleasmlib-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:

Table F-2 ORACLEASM Configure Prompts and Responses

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 Oracle Grid Infrastructure software owner user (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 ASMLIB driver file system is not a regular file system. It is used only by the Oracle ASM library to communicate with the Oracle ASM driver.

7. Enter the following command to load the `oracleasm` kernel module:

```
# /usr/sbin/oracleasm init
```

8. Repeat this procedure on all nodes in the cluster where you want to install Oracle RAC.

F.4.1.2.2 Configuring Disk Devices to Use Oracle ASM Library Driver on x86 Systems

To configure the disk devices to use in an Oracle ASM disk group, follow these steps:

1. If you intend to use IDE, SCSI, or RAID devices in the Oracle ASM disk group, then follow these steps:
 - a. If necessary, install or configure the shared 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, enter the following command:

```
# /sbin/fdisk -l
```

Depending on the type of disk, the device name can vary. [Table F-3](#) describes some types of disk paths:

Table F-3 Types of Linux Storage Disk Paths

Disk Type	Device Name Format	Description
IDE disk	/dev/hdxn	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sdxn	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.
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.

- 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 ASM disk:

```
# /usr/sbin/oracleasm createdisk DISK1 /dev/sdb1
```

In this example, `DISK1` is the name you assign to the disk.

Note: The disk names that you specify can contain uppercase letters, numbers, and the underscore character. They must start with an uppercase letter.

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.

F.4.1.2.3 Configuring Disk Devices to Use ASM Library Driver on IBM zSeries Systems 1.If you formatted the DASD with the compatible disk layout, then enter a command similar to the following to create a single whole-disk partition on the device:

```
# /sbin/fdasd -a /dev/dasdxxxx
```

2. Enter a command similar to the following to mark a disk as an ASM disk:

```
# /etc/init.d/oracleasm createdisk DISK1 /dev/dasdxxxx
```

In this example, DISK1 is a name that you want to assign to the disk.

Note: The disk names that you specify can contain uppercase letters, numbers, and the underscore character. They must start with an uppercase letter.

If you are using a multi-pathing disk driver with ASM, then make sure that you specify the correct logical device name for the disk.

3. To make the disk available on the other cluster nodes, enter the following command as root on each node:

```
# /etc/init.d/oracleasm scandisks
```

This command identifies shared disks attached to the node that are marked as ASM disks.

F.4.1.2.4 Administering the Oracle ASM Library Driver and Disks To administer the Oracle Automatic Storage Management library driver (ASMLIB) and disks, use the `/usr/sbin/oracleasm` initialization script with different options, as described in [Table F-4](#):

Table F-4 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 <i>labelname</i> is the name you want to use to mark the device, and <i>devicepath</i> is the path to the device: <code>oracleasm createdisk labelname devicepath</code>

Table F-4 (Cont.) Disk Management Tasks Using ORACLEASM

Task	Command Example	Description
Unmark a named disk device	oracleasm deletedisk VOL1	Use the <code>deletedisk</code> option to unmark a named disk device, where <i>diskname</i> is the name of the disk: oracleasm deletedisk <i>diskname</i> 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.
Determine if ASMLIB is using a disk device	oracleasm querydisk	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 <i>diskname_devicename</i> is the name of the disk or device that you want to query: oracleasm querydisk <i>diskname_devicename</i>
List Oracle ASMLIB disks	oracleasm listdisks	Use the <code>listdisks</code> option to list the disk names of marked ASMLIB disks.
Identify disks marked as ASMLIB disks	oracleasm scandisks	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	oracleasm renamedisk VOL1 VOL2	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 <i>manager</i> specifies the manager device, <i>label_device</i> specifies the disk you intend to rename, as specified either by OracleASM label name or by the device path, and <i>new_label</i> specifies the new label you want to use for the disk: oracleasm renamedisk [-l <i>manager</i>] [-v] <i>label_device</i> <i>new_label</i> 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.

F.4.1.3 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.

See Also: My Oracle Support site for updates to supported storage options:

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

F.4.1.3.1 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:

1. The initial path to the disk
2. The second path to the disk
3. 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.

F.4.1.3.2 Disk Scan Ordering 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_SCANORDER` variable specifies disks to be scanned first. 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.

F.4.1.3.3 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. For example:

```
multipaths {
    multipath {
        wwid                3600508b4000156d700012000000b0000
        alias                multipath
        ...
    }
    multipath {
        ...
        alias                mympath
        ...
    }
    ...
}
```

The default device name is in the format `/dev/mapper/mpath*` (or a similar path).

F.4.1.3.4 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.

F.4.1.4 Deinstalling Oracle ASMLIB

If you have Oracle ASMLIB installed but do not use it for storage persistence, you can deinstall it in rolling mode, one node at a time, as follows:

1. Login as `root`.
2. Stop Oracle ASM and any running database instance on the node:

```
srvctl stop asm -node node_name
srvctl stop instance -d db_unique_name -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 oracleasm library and tools RPMs:

```
rpm -e oracleasm-support
rpm -e oracleasm
```

5. Remove any oracleasm kernel driver RPMs provided by vendors:

```
rpm -e oracleasm
```

6. Check if any oracleasm RPMs remain:

```
rpm -qa | grep oracleasm
```

7. If any oracleasm configuration files remain, remove them:

```
rpm -qa | grep oracleasm | xargs rpm -e
```

Oracle ASMLIB and the associated RPMs are removed.

8. Start the Oracle Clusterware stack. Optionally, you can install and configure Oracle ASM Filter Driver (Oracle ASMFD) before starting the Oracle Clusterware stack.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about configuring storage device path persistence using Oracle ASM Filter Driver

F.4.2 Configuring Disk Devices Manually for Oracle ASM

This section contains the following information about preparing disk devices for use by Oracle ASM:

- [About Device File Names and Ownership for Linux](#)
- [Configuring a Permissions File for Disk Devices for Oracle ASM](#)

Note: The operation of udev depends on the Linux version, vendor, and storage configuration.

F.4.2.1 About Device File Names and Ownership for Linux

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 device file names and owners for voting files or Oracle Cluster Registry partitions, making them inaccessible when the server is restarted. For example, a voting file 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 ASMFD, then you do not need to ensure permissions and device path persistency in udev.

If you do not use ASMF, then you must create a custom rules file. 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, if you want to include the characteristic ACTION on the UDEV line, then 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.

F.4.2.2 Configuring a Permissions File for Disk Devices for Oracle ASM

To configure a permissions file for disk devices for Oracle ASM, complete the following tasks:

1. To obtain information about existing block devices, run the command `scsi_id (/sbin/scsi_id)` on storage devices from one cluster node to obtain their unique device identifiers. When running the `scsi_id` command with the `-s` argument, the device path and name passed should be that 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 of clusterware devices, 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
```

3. Using a text editor, create a UDEV rules file for the Oracle ASM devices, setting permissions to 0660 for the installation owner and the 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=="sd?1", OWNER="grid", GROUP="asmadmin", MODE="0660",
BUS=="scsi", PROGRAM=="/sbin/scsi_id", RESULT=="14f70656e66696c00000000"
KERNEL=="sd?2", OWNER="grid", GROUP="asmadmin", MODE="0660",
BUS=="scsi", PROGRAM=="/sbin/scsi_id", RESULT=="14f70656e66696c00000001"
KERNEL=="sd?3", OWNER="grid", GROUP="asmadmin", MODE="0660",
```

```
BUS=="scsi", PROGRAM=="sbin/scsi_id", RESULT=="14f70656e66696c00000002"
```

4. 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
```

5. Load updated block device partition tables on all member nodes of the cluster, using `/sbin/partprobe devicename`. You must do this as root.
6. 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:

```
# udevtest /block/sdb/sdb1
main: looking at device '/block/sdb/sdb1' 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, 'sdb1' becomes 'data1'
udev_device_event: device '/block/sdb/sdb1' 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/sdb1` to `/dev/data1`.

7. Enter the command to restart the UDEV service.

On Oracle Linux and Red Hat Enterprise Linux, the commands are:

```
# /sbin/udevcontrol reload_rules
# /sbin/start_udev
```

On SUSE Linux Enterprise Server, the command is:

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

F.5 Checking OCFS2 Version Manually

To check your OCFS2 version manually, enter the following commands:

```
modinfo ocfs2  
rpm -qa |grep ocfs2
```

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.

For information about OCFS2, refer to the following website:

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

Numerics

- 32-bit and 64-bit
 - software versions in the same cluster not supported, 2-2

A

- ACFS. *See* Oracle ACFS.
- ACFS-9427, A-16
- ACFS-9428, A-16
- aio-max-nr, F-6
- AMD 64
 - software requirements for, 4-10
- and Fixup script feature, 4-7
- architecture
 - checking system architecture, 2-1
- ASM
 - checking disk availability, 7-33
 - displaying attached disks
 - on Linux, 7-33
- ASM. *See* Oracle ASM
- ASMCA
 - Used to create disk groups for earlier Oracle Database releases on Oracle ASM, 9-7
- ASMFd
 - about, 7-27
- ASMLIB. *See* Oracle ASM library driver
- ASMSNMP, 1-4
- Automatic Storage Management Cluster File System.
See Oracle ACFS.
- Automatic Storage Management. *See* Oracle ASM.

B

- Bash shell
 - default user startup file, 6-21
 - setting shell limits, 6-23
- .bash_profile file, 6-21
- binaries
 - relinking, 9-8
- block devices
 - creating permissions file for Oracle Clusterware files, F-16
 - device name, F-11
 - unsupported for direct storage, 7-2

- BMC
 - configuring, 6-28
- BMC interface
 - preinstallation tasks, 6-25
- bonded addresses
 - must use same IP protocol, 5-3
- Bourne shell
 - default user startup file, 6-21
 - setting shell limits on Linux x86, 6-23

C

- C shell
 - default user startup file, 6-21
 - setting shell limits, 6-23
- candidate disks
 - marking, 7-27
- central inventory, 6-9
 - about, E-2
 - local on each cluster member node, 6-2
- central inventory. *See also* Oracle Inventory group
- changing host names, E-5
- checkdir error, 9-9, B-4
- checking disk availability for ASM, 7-33
- CHM. *See* Cluster Health Monitor
- chmod command, 7-19
- chown command, 7-19
- clients
 - connecting to SCANs, E-6
- cloning
 - cloning a Grid home to other nodes, 8-5
- cluster configuration file, 8-4
- cluster file system
 - storage option for data files, 7-5
- Cluster Health Monitor
 - enhancements for, xviii
 - repository in OCR location, B-16
- cluster name, 1-4
 - requirements for, 1-4
- cluster nodes
 - private network node interfaces, 1-4
 - private node names, E-8
 - public network node names and addresses, 1-4
 - public node names, 8-3
 - specifying uids and gids, 6-15
 - virtual node names, 1-4, E-6

- Cluster Time Synchronization Service, 4-29
- Cluster Verification Utility
 - cvuqdisk, 4-27
 - Fixup scripts, 4-7
 - user equivalency troubleshooting, A-8
- clusterware
 - requirements for third party clusterware, 1-4
- commands, 6-23
 - asmca, 7-29, 8-3, 9-3, B-12
 - asmcmd, 6-3
 - cat, 4-26
 - chmod, 7-19
 - chown, 7-19
 - crsctl, 8-7, 9-7, B-5, B-11
 - df, 2-2
 - env, 6-23
 - fdisk, 7-33, F-11, F-16
 - free, 2-1
 - groupadd, 6-15
 - id, 6-15
 - ipmitool, 6-28
 - lsdev, 7-33, F-11
 - lsmod, 6-27
 - mkdir, 7-19
 - modprobe, 6-27
 - nscd, 4-28
 - partprobe, F-18
 - passwd, 6-16
 - ping, 5-3
 - rootcrs.sh, 9-8
 - and deconfig option, 10-5
 - rootupgrade.sh, B-5
 - rpm, 4-27
 - rpm -qa | grep ssh, 4-3
 - sqlplus, 6-3
 - srvctl, B-5
 - swap, 2-1
 - swapon, 2-1
 - umask, 6-21
 - uname, 2-1, 4-27, F-8, F-9
 - unset, B-6
 - useradd, 6-5, 6-14, 6-16
 - xhost, 4-8
 - xterm, 4-8
 - yum, 3-1, F-8
- configuring kernel parameters, F-5
- cron jobs, 1-6, A-11
- CRS-5018 error, A-3
- CRS-5823 error, A-4
- ctsdd, 4-29
- custom database
 - failure groups for Oracle ASM, 7-24, 7-28
 - requirements when using Oracle ASM, 7-22
- Custom installation type
 - reasons for choosing, 6-10
- cvuqdisk, 4-27

D

data files

- creating separate directories for, 7-18, 7-19
- setting permissions on data file directories, 7-19
- storage options, 7-5
- data loss
 - minimizing with Oracle ASM, 7-24, 7-28
- database files
 - supported storage options, 7-2
- databases
 - Oracle ASM requirements, 7-22
- dba group. *See* OSDBA group
- DBCA
 - no longer used for Oracle ASM disk group administration, 9-7
 - Used to create server pools for earlier Oracle Database releases, 9-6
- dbca.rsp file, C-3
- default file mode creation mask
 - setting, 6-21
- default Linux installation
 - recommendation for, 4-3
- deinstallation, 10-1
- Deinstallation tool
 - about, 10-6
 - example, 10-9
 - previous Grid home, 10-9
 - Restriction for Oracle Flex Clusters and -lastnode flag, 10-5
 - roothas.sh, 10-6
- deinstalling previous Grid home, 10-9
- device names
 - IDE disks, F-11
 - RAID disks, F-11
 - SCSI disks, F-11
- df command, 2-2, 6-23
- DHCP
 - and GNS, 5-5
- Direct NFS Client
 - disabling, 7-20
 - enabling, 7-15, 7-16
 - for data files, 7-10
 - minimum write size value for, 7-11
- directory
 - creating separate data file directories, 7-18, 7-19
 - permission for data file directories, 7-19
- disk group
 - Oracle ASM, 7-21
 - recommendations for Oracle ASM disk groups, 7-21
- disk groups
 - recommendations for, 7-21
- disk space
 - checking, 2-1
 - requirements for preconfigured database in Oracle ASM, 7-22
- disks
 - checking availability for ASM, 7-33
 - checking availability for Oracle ASM, F-11
 - displaying attached disks, F-11
 - on Linux, 7-33
- disks. *See also* Oracle ASM disks

DISPLAY environment variable
setting, 6-22

E

emulator
installing from X emulator, 4-8
enterprise.rsp file, C-3
env command, 6-23
environment
checking settings, 6-23
configuring for Oracle user, 6-20
environment variables
DISPLAY, 6-22
ORACLE_BASE, 6-22
ORACLE_HOME, 6-3, 6-22, B-6
ORACLE_SID, 6-22, B-6
removing from shell startup file, 6-22
SHELL, 6-21
TEMP and TMPDIR, 2-2, 6-22
ephemeral ports
setting manually, F-7
errors
X11 forwarding, 6-24, F-5
errors using Opatch, 9-9, B-4
Exadata
and rp_filter, 5-17
relinking binaries example for, 9-8
examples
Oracle ASM failure groups, 7-24, 7-28

F

failure group
characteristics of Oracle ASM failure group, 7-24, 7-28
examples of Oracle ASM failure groups, 7-24, 7-28
Oracle ASM, 7-21
fdisk command, 7-33, F-11
fencing
and IPMI, 1-2, 6-25
file mode creation mask
setting, 6-21
file system
storage option for data files, 7-5
file systems, 7-7
file-max, F-6
files
.bash_profile, 6-21
dbca.rsp, C-3
editing shell startup file, 6-21
enterprise.rsp, C-3
.login, 6-21
oraInst.loc, 6-2
.profile, 6-21
response files, C-3
filesets, 4-9
Fixup script, 4-7

G

GFS, 7-7
gid
identifying existing, 6-15
specifying, 6-15
specifying on other nodes, 6-15
GID changes for existing install owners
unsupported, 6-4
globalization
support for, 1-6
GNS
about, 5-6
configuring, 5-5
GNS client clusters
and GNS client data file, 5-7
GNS client data file required for installation, 5-6
name resolution for, 5-6
GNS client data file
how to create, 5-7
GNS virtual IP address, 1-4
GPFS, 7-7
Grid home
and Oracle base restriction, 6-7
disk space for, 1-2, 2-3
minimum required space for, 2-3
unlocking, 9-8
Grid Infrastructure Management Repository, xviii
shared disk file sizes, 7-8
grid naming service. *See* GNS
Grid user, 6-9
Oracle base for, 2-3
grid_install.rsp file, C-3
group IDs
identifying existing, 6-15
specifying, 6-15
specifying on other nodes, 6-15
groups
checking for existing OINSTALL group, 6-1
creating identical groups on other nodes, 6-15
creating the Oracle ASM group, 6-13
creating the OSDBA for ASM group, 6-13
creating the OSDBA group, 6-12
OINSTALL, 6-2
OSASM (asmadmin), 6-11
OSBACKUPDBA (backupdba), 6-11
OSDBA (dba), 6-10
OSDBA for ASM (asmdba), 6-11
OSDBA group (dba), 6-10
OSDGDBA (dgdba), 6-11
OSKMDBA (kmdba), 6-11
OSOPER (oper), 6-10
OSOPER for ASM (asmoper), 6-12
required for Oracle Software Owner users, 6-9
specifying when creating users, 6-15
using NIS, 6-9, 6-15

H

high availability IP addresses, 5-3
host names

- changing, E-5
- legal host names, 1-4

Hub Nodes, 5-9, E-8

HugePages, D-1

I

id command, 6-15

IDE disks

- device names, F-11

inaccessible nodes

- upgrading, B-11

INS-32026 error, 6-7

installation

- and cron jobs, 1-6
- and globalization, 1-6
- cloning a Grid infrastructure installation to other nodes, 8-5
- failed, A-18
- interrupted, A-18
- response files, C-3
 - preparing, C-3, C-4
 - templates, C-3
- silent mode, C-5
- using cluster configuration file, 8-4

installation types

- and Oracle ASM, 7-22

installer screen

- Failure Isolation Support, 6-25
- Node Selection, A-6
- Summary, C-4

interconnect, 1-4

interfaces, 1-4

- requirements for private interconnect, E-6

intermittent hangs

- and socket files, 8-7

IP protocol

- and redundant interfaces, 5-3

IP_LOCAL_PORT_RANGE

- setting manually, F-7

ip_local_port_range, F-6

IPMI

- addresses not configurable by GNS, 6-26
- configuring driver for, 6-26
- preinstallation tasks, 6-25
- preparing for installation, 1-2

IPv4 requirements, 5-3

IPv6 requirements, 5-3

IPv6 support

- about, xix

J

JDK requirements, 4-9

job role separation users, 6-9

K

kernel parameters

- configuring, F-5

kernel requirements

Linux x86-64, 4-18, 4-19, 4-20

Korn shell

- default user startup file, 6-21
- setting shell limits, 6-23

ksh. *See* Korn shell

L

Leaf Nodes, 5-9, E-8

legal host names, 1-4

Linux

- cvuqdisk package, 4-27
- determining distribution of, 4-26
- displaying attached disks, 7-33
- rp_filter setting for multiple interconnects, 5-17

Linux x86-64

- software requirements, 4-18, 4-19, 4-20
- software requirements for, 4-10

log file

- how to access during installation, 8-3

.login file, 6-21

lsdev command, 7-33, F-11

LVM

- recommendations for Oracle ASM, 7-21

M

mask

- setting default file mode creation mask, 6-21

minimal Linux installation

- recommendation for, 4-2

mixed binaries, 4-10

mkdir command, 7-19

mode

- setting default file mode creation mask, 6-21

multiple interconnects

- rp_filter setting for, 5-17

multiple Oracle homes, 6-3, 7-19

My Oracle Support, 9-1

N

Name Service Cache Daemon

- enabling, 4-28

Net Configuration Assistant (NetCA)

- response files, C-6
- running at command prompt, C-6

netca, 8-3

netca.rsp file, C-3

Network Information Services. *See* NIS

network port ranges, F-7

networks

- for Oracle Flex Clusters, 5-9, E-8
- IP protocol requirements for, 5-3
- Oracle Flex ASM, 1-4, 5-10, E-8

NFS, 7-7, 7-13

- and data files, 7-12
- and Oracle Clusterware files, 7-7
- buffer size parameters for, 7-13, 7-14
- Direct NFS Client, 7-10
- for data files, 7-12

rsize, 7-13
NIS
alternative to local users and groups, 6-9
noninteractive mode. *See* response file mode

O

OCFS2, 7-7
checking version of, F-19
OCR. *See* Oracle Cluster Registry
ocrconfig error, A-5
OINSTALL group
about, E-2
and oraInst.loc, 6-2
checking for existing, 6-1
creating on other nodes, 6-15
creating the oraInventory group, 6-2
system privileges granted by, 6-2, 6-4
OINSTALL group. *See also* Oracle Inventory group
Opatch, 9-9, B-4
OpenSSH
and Oracle Preinstallation RPM, 4-3, 4-4, 4-11, 4-12, 4-14, 4-15, 4-16
command to check to see if installed, 4-3
operating system
different on cluster members, 4-10
limitation for Oracle ACFS, E-9
missing packages, A-10
requirements, 4-9
operating system authentication for system
privileges, 6-10
operating system privileges groups, 6-10
optimal flexible architecture
and oraInventory directory, E-3
Oracle ACFS
about, E-9
Installing Oracle RAC binaries not supported on
Oracle Flex Cluster, 7-5
Oracle ASM
block device names, F-11
candidate disks, 7-27
characteristics of failure groups, 7-24, 7-28
checking disk availability, F-11
configuring disks for Oracle ASM, 7-27
Configuring the Oracle ASM library driver, F-9
creating the OSDBA for ASM group, 6-13
disk groups, 7-21
displaying attached disks, F-11
failure groups, 7-21
examples, 7-24, 7-28
identifying, 7-24, 7-28
identifying available disks, F-11
installation as part of Oracle Grid Infrastructure
installation, xxi
Oracle ASM library driver (ASMLIB), 8-3
OSASM for ASM administrator, 6-11
OSDBA for ASM group, 6-11
OSOPER for ASM group, 6-12
recommendations for disk groups, 7-21
space required for Oracle Clusterware files, 7-22
space required for preconfigured database, 7-22
storage option for data files, 7-5
storing Oracle Clusterware files on, 7-2
Oracle ASM clients, 5-10
Oracle ASM disks
marking, 8-3
Oracle ASM Filter Driver
about, 7-27
Oracle ASM group
creating, 6-13
Oracle ASM library driver
configuring system startup options for, F-14
How to configure, F-9
Oracle ASM library driver (oracleasm)
installing, F-8
Oracle base directory
Grid home must not be in an Oracle Database
Oracle base, 6-7
Grid homes not permitted under, E-2
minimum disk space for, 1-2, 2-3
requirements for Oracle Grid Infrastructure, E-4
Oracle base for Oracle Grid Infrastructure installation
owner (Grid user)
minimum required space for, 2-3
Oracle Cluster Registry
configuration of, 1-6
mirroring, 7-8
partition sizes, 7-8
permissions file to own block device
partitions, F-16
supported storage options, 7-2
Oracle Clusterware
and file systems, 7-7
installation as part of Oracle Grid Infrastructure
installation, xxi
installing, 8-1
supported storage options for, 7-2
upgrading, 7-8
Oracle Clusterware files
Oracle ASM disk space requirements, 7-22
Oracle Clusterware Installation Guide
replaced by Oracle Grid Infrastructure Installation
Guide, 8-1
Oracle Database
creating data file directories, 7-18, 7-19
data file storage options, 7-5
privileged groups, 6-10
requirements with Oracle ASM, 7-22
Oracle Database Configuration Assistant
response file, C-3
Oracle Disk Manager
and Direct NFS Client, 7-15
Oracle Flex ASM
about, E-8
and Oracle ASM clients, 1-4, 5-10, E-8
networks, 1-4, 5-10, E-8
Oracle Flex Clusters
about, xviii
and Hub Nodes, 9-6, E-8
and Leaf Nodes, 9-6

- and Oracle Flex ASM, 5-9, E-8
- restrictions for Oracle ACFS, 7-5
- Oracle Grid Infrastructure owner (grid), 6-9
- Oracle Grid Infrastructure response file, C-3
- Oracle home
 - and asmcmd errors, 6-3
 - ASCII path restriction for, 1-3
 - multiple Oracle homes, 6-3, 7-19
- Oracle Inventory group
 - about, E-2
 - checking for existing, 6-1
 - creating, 6-2
 - creating on other nodes, 6-15
 - oraInst.loc file and, 6-2
- Oracle Linux
 - and Oracle Preinstallation RPM
 - accounts configured by, 4-5
- Oracle Net Configuration Assistant
 - response file, C-3
- Oracle patch updates, 9-1
- Oracle Preinstallation RPM
 - about, 4-4
 - included with Oracle Linux, 4-5
 - installing, 3-2
- Oracle Real Application Clusters
 - configuring disks for Oracle ASM, F-10
- Oracle Restart
 - downgrading, A-15
- Oracle Software Owner users, F-10
 - configuring environment for, 6-20
 - creating, 6-3, 6-4, 6-13, 6-14
 - creating on other nodes, 6-15
 - description, 6-9
 - determining default shell, 6-21
 - required group membership, 6-9
 - setting shell limits for, 6-23
- Oracle Universal Installer
 - response files
 - list of, C-3
- Oracle Upgrade Companion, 4-6
- Oracle user
 - and Oracle Preinstallation RPM, 4-5
 - configuring environment for, 6-20
 - creating, 6-3, 6-4, 6-13, 6-14
 - creating on other nodes, 6-15
 - description, 6-9
 - determining default shell, 6-21
 - required group membership, 6-9
 - setting shell limits for, 6-23
- Oracle user. *See also* Oracle Software Owner users
- Oracle Validated RPM. *See* Oracle Preinstallation RPM
- ORACLE_BASE environment variable
 - removing from shell startup file, 6-22
- ORACLE_HOME environment variable
 - removing from shell startup file, 6-22
- ORACLE_SID environment variable
 - removing from shell startup file, 6-22
- oracleasm RPM
 - installing, F-8

- OraclePreinstallation RPM
 - and OpenSSH requirement, 4-4
- oraInst.loc
 - and central inventory, 6-2
 - contents of, 6-2
- oraInst.loc file
 - location, 6-2
 - location of, 6-2
- oraInventory, 6-9
 - about, E-2
 - creating, 6-2
- oraInventory. *See also* Oracle Inventory group
- OSASM group, 6-11
 - about, 6-11
 - creating, 6-13
- OSBACKUPDBA group (backupdba), 6-11
- OSDBA for ASM group, 6-11
 - about, 6-11
- OSDBA group
 - and SYSDBA system privileges, 6-10
 - creating, 6-12
 - creating on other nodes, 6-15
 - description, 6-10
- OSDBA group for ASM
 - creating, 6-13
- OSDGDBA group (dgdbs), 6-11
- OSKMDBA group (kmdbs), 6-11
- OSOPER for ASM group, 6-12
 - about, 6-11
 - creating, 6-13
- OSOPER group
 - and SYSOPER system privileges, 6-10
 - creating, 6-12
 - creating on other nodes, 6-15
 - description, 6-10
- OUI, 4-7

P

- package cvuqdisk not installed, 4-27
- package requirements
 - Linux x86-64, 4-18, 4-19, 4-20
- packages
 - checking on Linux, 4-27
 - default Linux installation and, 4-10
- partition
 - using with Oracle ASM, 7-21
- passwd command, 6-16
- patch updates
 - download, 9-1
 - install, 9-1
 - My Oracle Support, 9-1
- permissions
 - for data file directories, 7-19
- policy-managed databases
 - and SCANS, E-7
- postinstallation
 - patch download and install, 9-1
- preconfigured database
 - Oracle ASM disk space requirements, 7-22

- requirements when using Oracle ASM, 7-22
- primary host name, 1-4
- privileged groups
 - for Oracle Database, 6-10
- processor
 - checking system architecture, 2-1
 - /proc/sys/fs/aio-max-nr, F-6
 - /proc/sys/fs/file-max file, F-6
 - /proc/sys/kernel/sem file, F-6
 - /proc/sys/kernel/shmall, F-6
 - /proc/sys/kernel/shmall file, F-6
 - /proc/sys/kernel/shmmax, F-6
 - /proc/sys/kernel/shmmni file, F-6
 - /proc/sys/kernel/shmni, F-6
 - /proc/sys/net/core/rmem_default file, F-6
 - /proc/sys/net/core/rmem_max file, F-6
 - /proc/sys/net/core/wmem_default file, F-6
 - /proc/sys/net/core/wmem_max file, F-6
 - /proc/sys/net/ipv4/ip_local_port_range, F-6
- .profile file, 6-21
- PRVF-5150 error, A-3
- public node name
 - and primary host name, 1-4

R

- RAC
 - configuring disks for ASM on Linux, 7-33
- RAID
 - and mirroring Oracle Cluster Registry and voting files, 7-8
 - recommended Oracle ASM redundancy level, 7-21
- RAID disks
 - device names, F-11
- raw devices
 - unsupported, 7-2
 - upgrading existing partitions, 7-8
- recommendations
 - client access to the cluster, 9-4
- recovery files
 - supported storage options, 7-2
- Red Hat Package Manager. *See* RPM
- redundancy level
 - and space requirements for preconfigured database, 7-22
- Redundant Interconnect Usage, 5-3
- redundant interfaces
 - must use same IP protocol, 5-3
- relinking Oracle Grid Infrastructure home binaries, 9-8, 10-3, 10-4
- requirements, 7-22
- response file installation
 - preparing, C-3
 - response files
 - templates, C-3
 - silent mode, C-5
- response file mode
 - about, C-1
 - reasons for using, C-2

- See also* response files, silent mode, C-1
- response files
 - about, C-1
 - creating with template, C-3
 - dbca.rsp, C-3
 - enterprise.rsp, C-3
 - general procedure, C-2
 - grid_install.rsp, C-3
 - Net Configuration Assistant, C-6
 - netca.rsp, C-3
 - passing values at command line, C-1
 - specifying with Oracle Universal Installer, C-5
- response files. *See also* silent mode
- rmem_default, F-6
- rmem_max, F-6
- root user
 - logging in as, 4-8
- rootcrs.sh
 - restriction for Oracle Flex Cluster
 - deinstallation, 10-5
- roothas.sh, 10-6
- root.sh, 8-3
- RPMs
 - checking, 4-27
 - default Linux installation and, 4-10
- rsh
 - no longer supported for inter-node communication, 4-30
- rsize parameter, 7-13
- run level, 2-2

S

- SCAN address, 1-4
- SCAN listener, E-7
- SCANs, 1-4, 5-8
 - client access, 9-4
 - configuring, 1-4
 - description, 9-4
 - understanding, E-6
 - use of SCANs required for clients of policy-managed databases, E-7
- SCSI disks
 - device names, F-11
- security
 - dividing ownership of Oracle software, 6-8
- sem file, F-6
- semnmi parameter
 - recommended value on Linux, F-6
- semnms parameter
 - recommended value on Linux, F-6
- semmsl parameter
 - recommended value on Linux, F-6
- semopm parameter
 - recommended value on Linux, F-6
- Server Parameter File, A-15
- setting shell limits, 6-23
- shell
 - determining default shell for Oracle user, 6-21
- SHELL environment variable

- checking value of, 6-21
- shell limits
 - setting on Linux, 6-23
- shell startup file
 - editing, 6-21
 - removing environment variables, 6-22
- shmall, F-6
- shmmax, F-6
- shmmni, F-6
- silent mode
 - about, C-1
 - reasons for using, C-2
 - See also* response files., C-1
- silent mode installation, C-5
- single client access names. *See* SCANS
- software requirements, 4-9
 - checking software requirements, 4-26
 - for Linux x86-64, 4-18, 4-19, 4-20
- SPF file, A-15
- ssh
 - and X11 Forwarding, 6-24
 - automatic configuration from OUI, 4-30
 - configuring, F-1
 - supported version of, F-1
 - when used, 4-30
- Standard cluster
 - upgrades result in, B-4
- startup file
 - for shell, 6-21
- storage
 - marking Oracle ASM candidate disks, 8-3
- stty
 - suppressing to prevent installation errors, 6-25
- supported storage options
 - Oracle Clusterware, 7-2
- suppressed mode
 - reasons for using, C-2
- SYSASM system privileges, 6-11
- SYSBACKUPDBA system privileges, 6-11
- SYSDBA for ASM system privileges, 6-11
- SYSDBA system privileges
 - associated group, 6-10
- SYSDGDBA system privileges, 6-11
- SYSKMDBA system privileges, 6-11
- SYSOPER for ASM system privileges, 6-12
- SYSOPER system privileges
 - associated group, 6-10
- system architecture
 - checking, 2-1
- system privileges
 - SYSASM, 6-11
 - SYSBACKUPDBA, 6-11
 - SYSDBA, 6-10
 - SYSDBA for ASM, 6-11
 - SYSDGDBA, 6-11
 - SYSKMDBA, 6-11
 - SYSOPER, 6-10
 - SYSOPER for ASM, 6-12

T

- TCP ephemeral ports, F-7
- tcsh shell
 - setting shell limits, 6-23
- TEMP environment variable, 2-2
 - setting, 6-22
- temporary directory, 2-1
- temporary disk space
 - checking, 2-1
 - freeing, 2-1
- terminal output commands
 - suppressing for Oracle installation owner
 - accounts, 6-25
- third party multicast DNS
 - restrictions, 5-6
- /tmp directory
 - checking space in, 2-1
 - freeing space in, 2-1
- TMPDIR environment variable, 2-2
 - setting, 6-22
- Transparent HugePages
 - disabling, D-2
- Troubleshooting
 - DBCA does not recognize Oracle ASM disk size
 - and fails to create disk groups, 9-7
- troubleshooting
 - ACFS-9427, A-16
 - ACFS-9428, A-16
 - and deinstalling, 10-1
 - asmcmd errors and Oracle home, 6-3
 - automatic SSH configuration from OUI, 4-31, 6-3
 - CRS 1137, A-17
 - CRS-0219, A-15
 - CRS-1604, A-10
 - CRS-2529, A-15
 - disk space errors, 1-3
 - DISPLAY errors, 6-24
 - environment path errors, 1-3
 - error messages, A-2
 - Fatal Timeout before authentication, A-6
 - garbage strings in script inputs found in log
 - files, 6-25
 - GID, 6-4
 - installation owner environment variables and
 - installation errors, B-6
 - intermittent hangs, 8-7
 - log file, 8-3
 - missing operating system packages, A-10
 - nfs mounts, 4-28
 - permissions errors and oraInventory, E-2
 - PROT-8 error, A-6
 - PRVE-0038 error, A-6
 - public network failures, 4-28
 - required secure shell (OpenSSH), 4-3, 4-4, 4-11,
 - 4-12, 4-14, 4-15, 4-16
 - root.sh errors, 10-5
 - run level error, 2-2
 - SQLPlus errors and Oracle home, 6-3
 - ssh, F-2
 - ssh configuration failure, F-2

- ssh errors, 6-25
- SSH timeouts, A-6
- stty errors, 6-25
- UID, 6-4
- Unable to clear device, A-14
- Unable to find candidate disks, A-14
- Unable to open ASMLib, A-14
- unconfiguring Oracle Clusterware to fix causes of
 - root.sh errors, 10-5
- unexplained installation errors, 1-6, A-11
- user equivalency, A-8, F-2
- user equivalency error due to different user or
 - group IDs, 6-4, 6-14
- user equivalency errors, 6-3, E-3
- user equivalency issues, 6-4
- X11 forwarding error, 6-24

U

- UDP ephemeral ports, F-7
- uid
 - identifying existing, 6-15
 - specifying, 6-15
 - specifying on other nodes, 6-15
- UID changes for existing install owners
 - unsupported, 6-4
- umask, 6-23
- umask command, 6-21, 6-23
- Unbreakable Enterprise Kernel for Linux
 - about, 4-3
 - and `rp_filter` setting, 5-17
- Unbreakable Linux Kernel. *See* Unbreakable Enterprise Kernel for Linux.
- unconfiguring Oracle Clusterware, 10-5
- uninstall, 10-1
- uninstalling, 10-1
- unreachable nodes
 - upgrading, B-11
- unset installation owners environment variables, B-6
- upgrade
 - failed or incomplete, A-17
- upgrades
 - and Leaf Nodes, B-4
 - and OCR partition sizes, 7-8
 - and SCANS, E-7
 - and voting file sizes, 7-8
 - best practices, 4-6
 - of Oracle ASM, B-12
 - restrictions for, B-3
 - shared Oracle Clusterware home to local Grid homes, E-2
 - unseting environment variables for, B-6
- upgrading
 - inaccessible nodes, B-11
- user equivalence
 - testing, A-8
- user equivalency errors
 - groups and users, 6-4, 6-14
- user IDs
 - identifying existing, 6-15

- specifying, 6-15
- specifying on other nodes, 6-15

useradd command, 6-5, 6-14, 6-16

users

- creating identical users on other nodes, 6-15
- creating the Grid user, 6-3
- creating the Oracle user, 6-4, 6-13, 6-14
- Oracle Software Owner users, 6-9
- setting shell limits for, 6-23
- setting shell limits for users on Linux, 6-23
- specifying groups when creating, 6-15
- using NIS, 6-9, 6-15

V

- vendor clusterware
 - and cluster names for Oracle Grid Infrastructure, 1-4
- voting files
 - configuration of, 1-6
 - mirroring, 7-8
 - partition sizes, 7-8
 - supported storage options, 7-2

W

- wmem_default, F-6
- wmem_max, F-6
- workstation
 - installing from, 4-8
- wsize parameter, 7-13
- wtmax, 7-11
 - minimum value for Direct NFS Client, 7-11

X

- X emulator
 - installing from, 4-8
- X terminal
 - installing from, 4-8
- X Window System
 - enabling remote hosts, 4-8
- X11 forwarding errors, 6-24, F-5
- xhost command, 4-8
- xterm command, 4-8
- xtitle
 - suppressing to prevent installation errors, 6-25

Y

- Yum repository, 3-1

